

BTeV Trigger (WBS 1.8)
and
Data Acquisition System (WBS 1.9)

Erik Gottschalk (WBS 1.8)

Klaus Honscheid, Margaret Votava (WBS 1.9)

- Introduction and overview of the BTeV trigger & data acquisition system (DAQ)
 - **WBS 1.8** – Trigger electronics & software
 - Project description
 - Project organization
 - Costs
 - Schedule
 - Milestones
 - Risk assessment
 - Response to CD-1 recommendations
 - **WBS 1.9** – DAQ electronics & software
 - Project description
 - ...
 - Presentations prepared for the breakout sessions
-

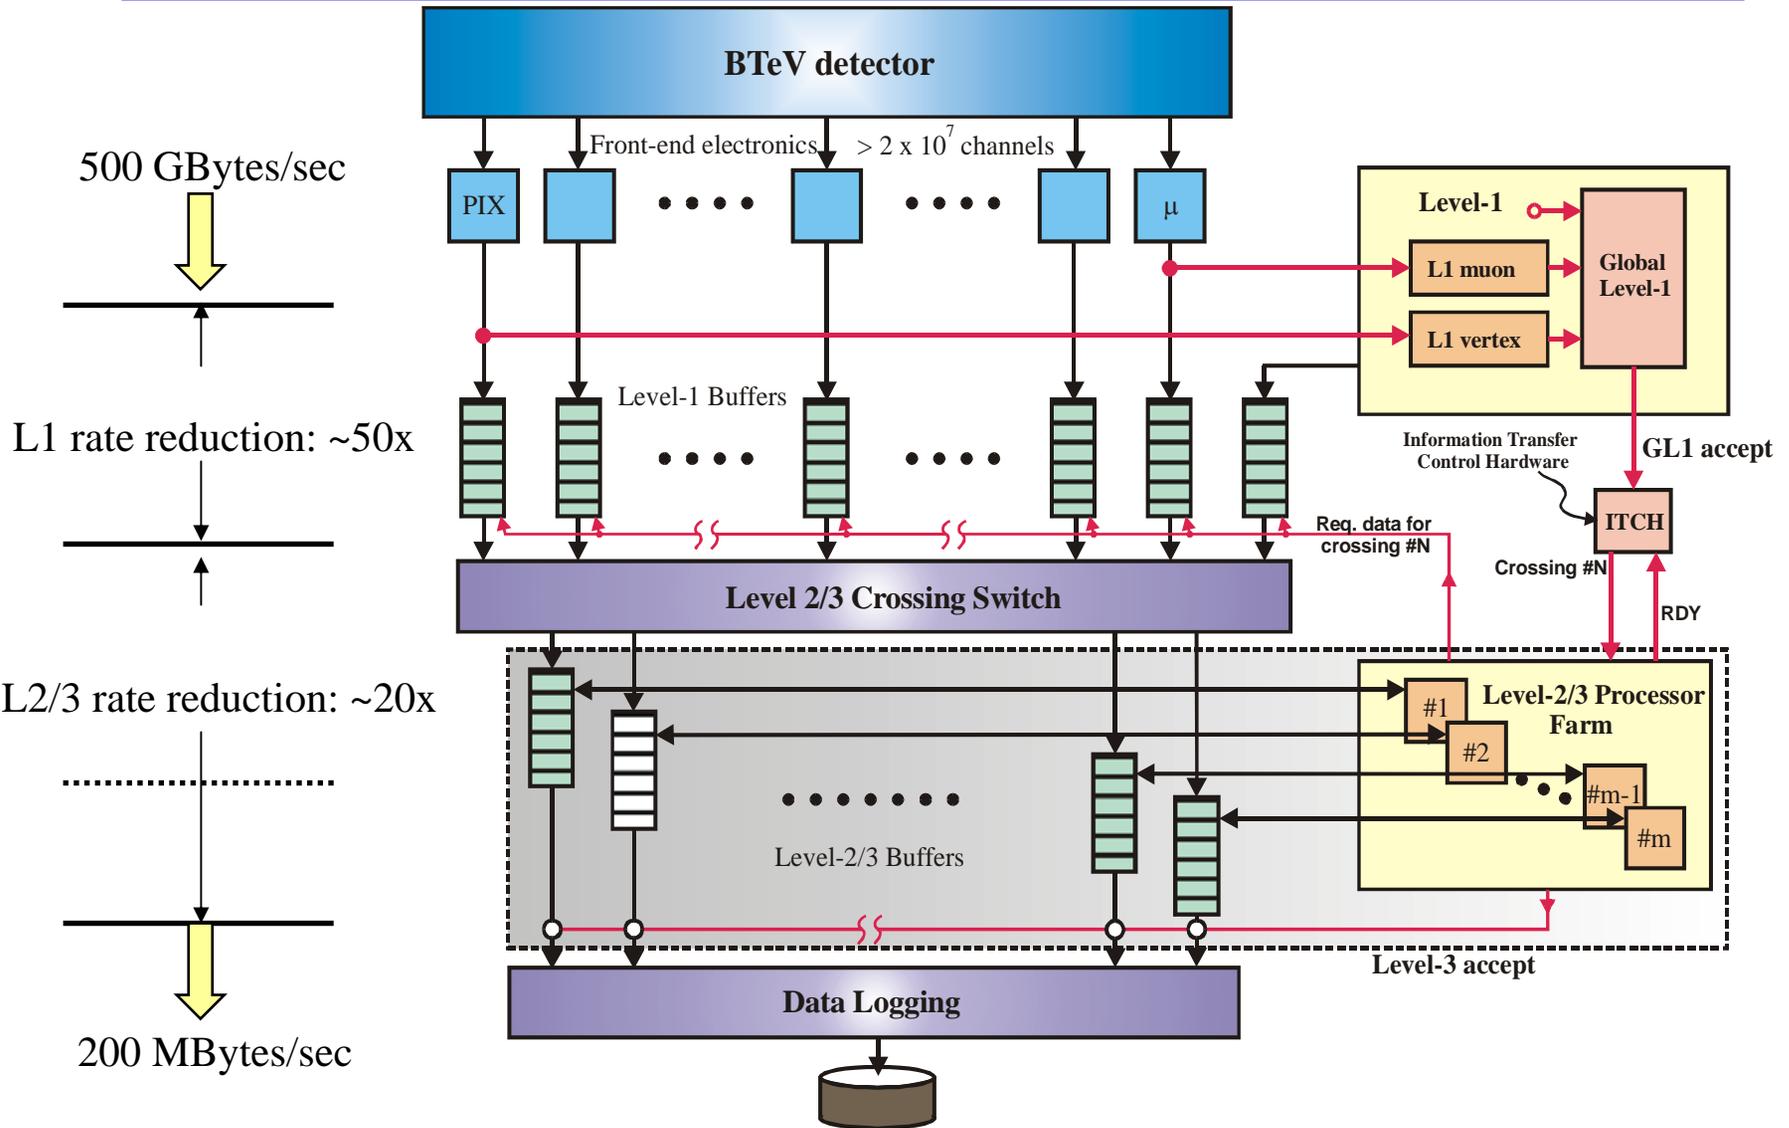
- The challenge for the BTeV trigger and data acquisition system is to reconstruct particle tracks and interaction vertices for **EVERY** interaction that occurs in the BTeV detector, and to select interactions with *B* decays.
- The trigger performs this task using 3 stages, referred to as Levels 1, 2, and 3:
 - “L1” – looks at every interaction, and rejects at least 98% of min. bias background
 - “L2” – uses L1 computed results & performs more refined analyses for data selection
 - “L3” – performs a complete analysis using all of the data for an interaction

Reject > 99.9% of background. Keep > 50% of *B* events.

- The data acquisition system saves all of the data in memory for as long as necessary to analyze each interaction (~ 1 millisecond on average for L1), and moves data to L2/3 processing units and archival data storage for selected interactions.
- The key ingredients that make it possible to meet this challenge:
 - BTeV pixel detector with its exceptional pattern recognition capabilities
 - Rapid development in technology – FPGAs, processors, networking

Note: see glossary at the end of this talk

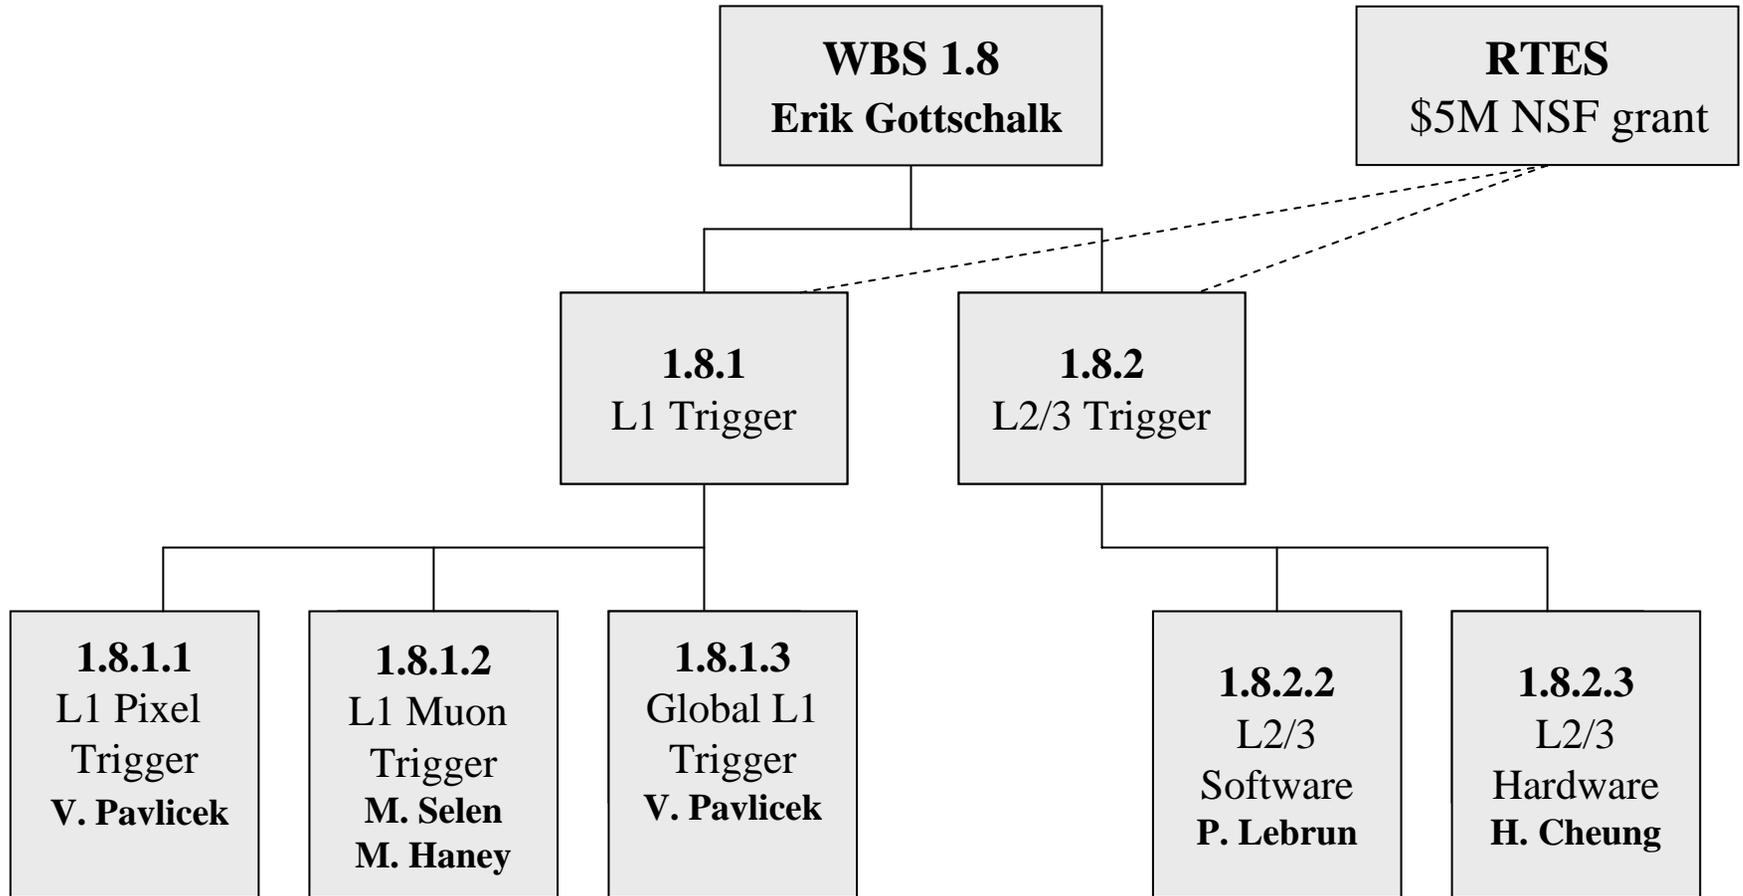
Block Diagram of the Trigger & DAQ



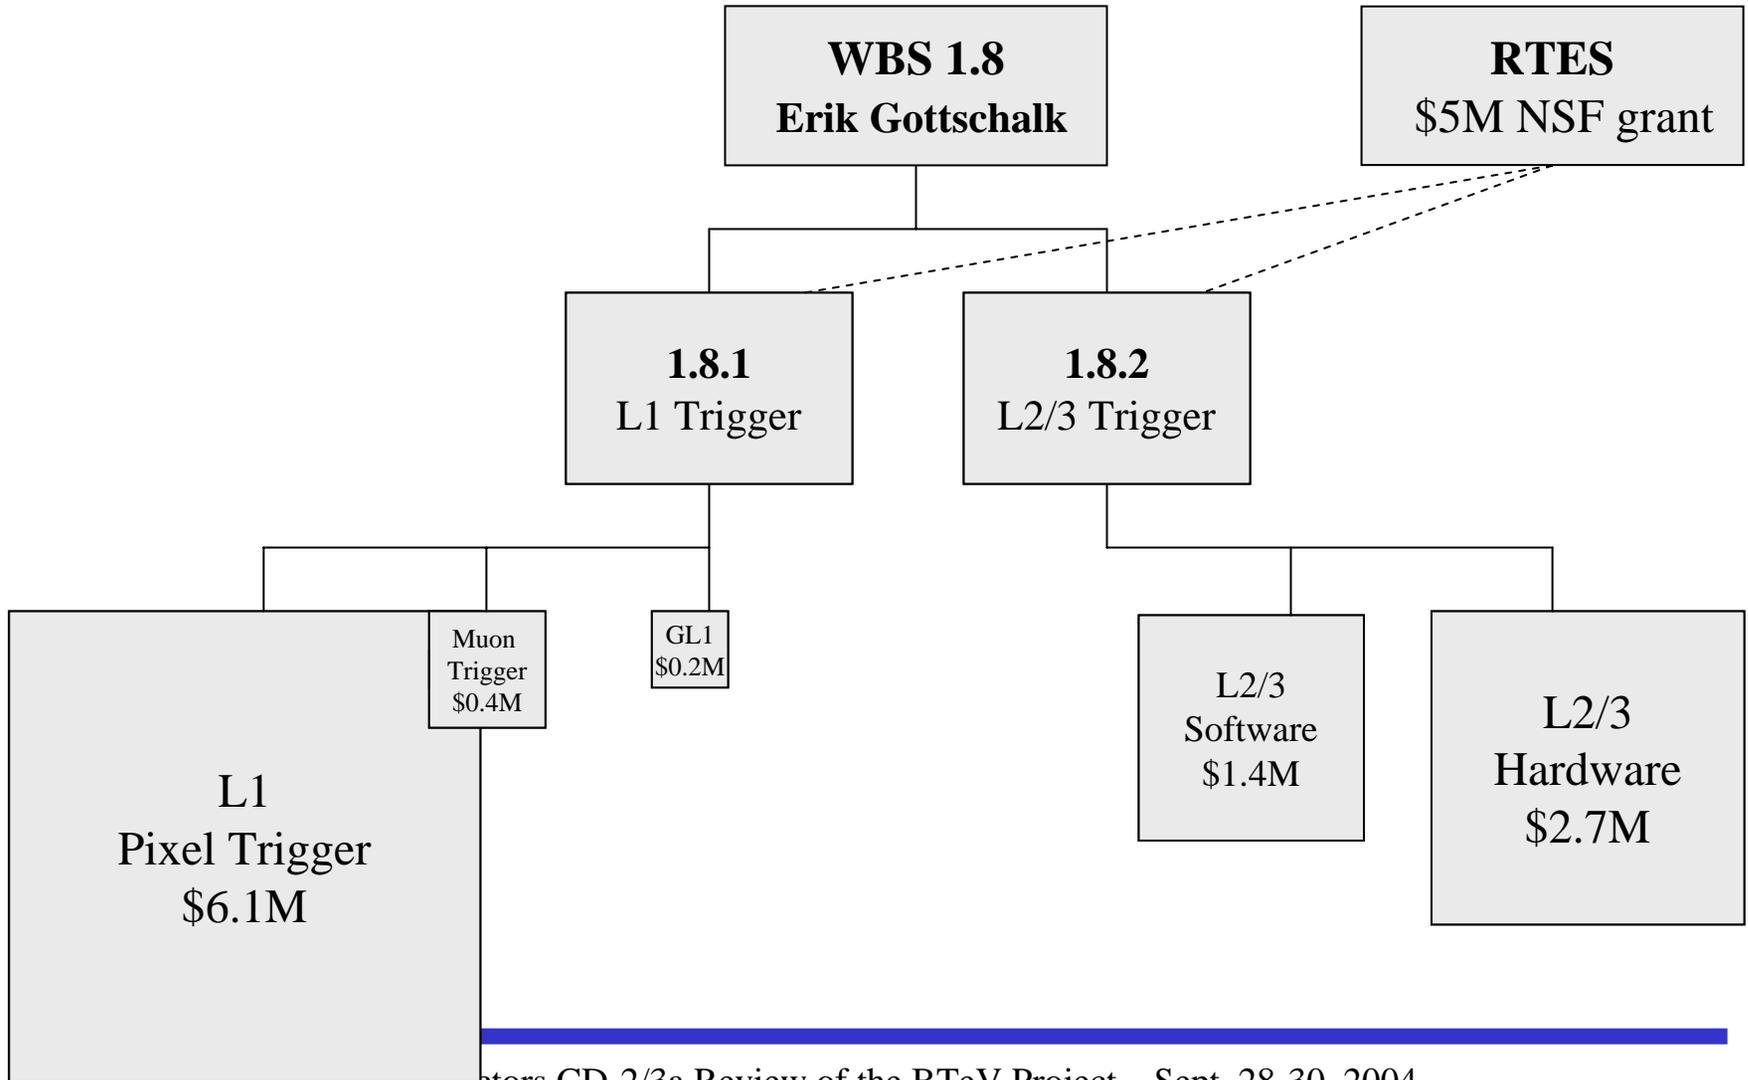
- L1 pixel trigger (FPGAs, L1 Switch, L1 Farm)
- L1 muon trigger (same hardware as L1 pixel trigger)
- Global Level 1 trigger (same processing hardware)
- L2/3 hardware (Linux PC farm)
- L2/3 software (similar to HEP “offline” analysis)
- RTES software (fault detection and mitigation)

Base cost: \$11.1M (Material: \$7.9M, Labor: \$3.3M)
+ \$5M grant for RTES (NSF ITR program)

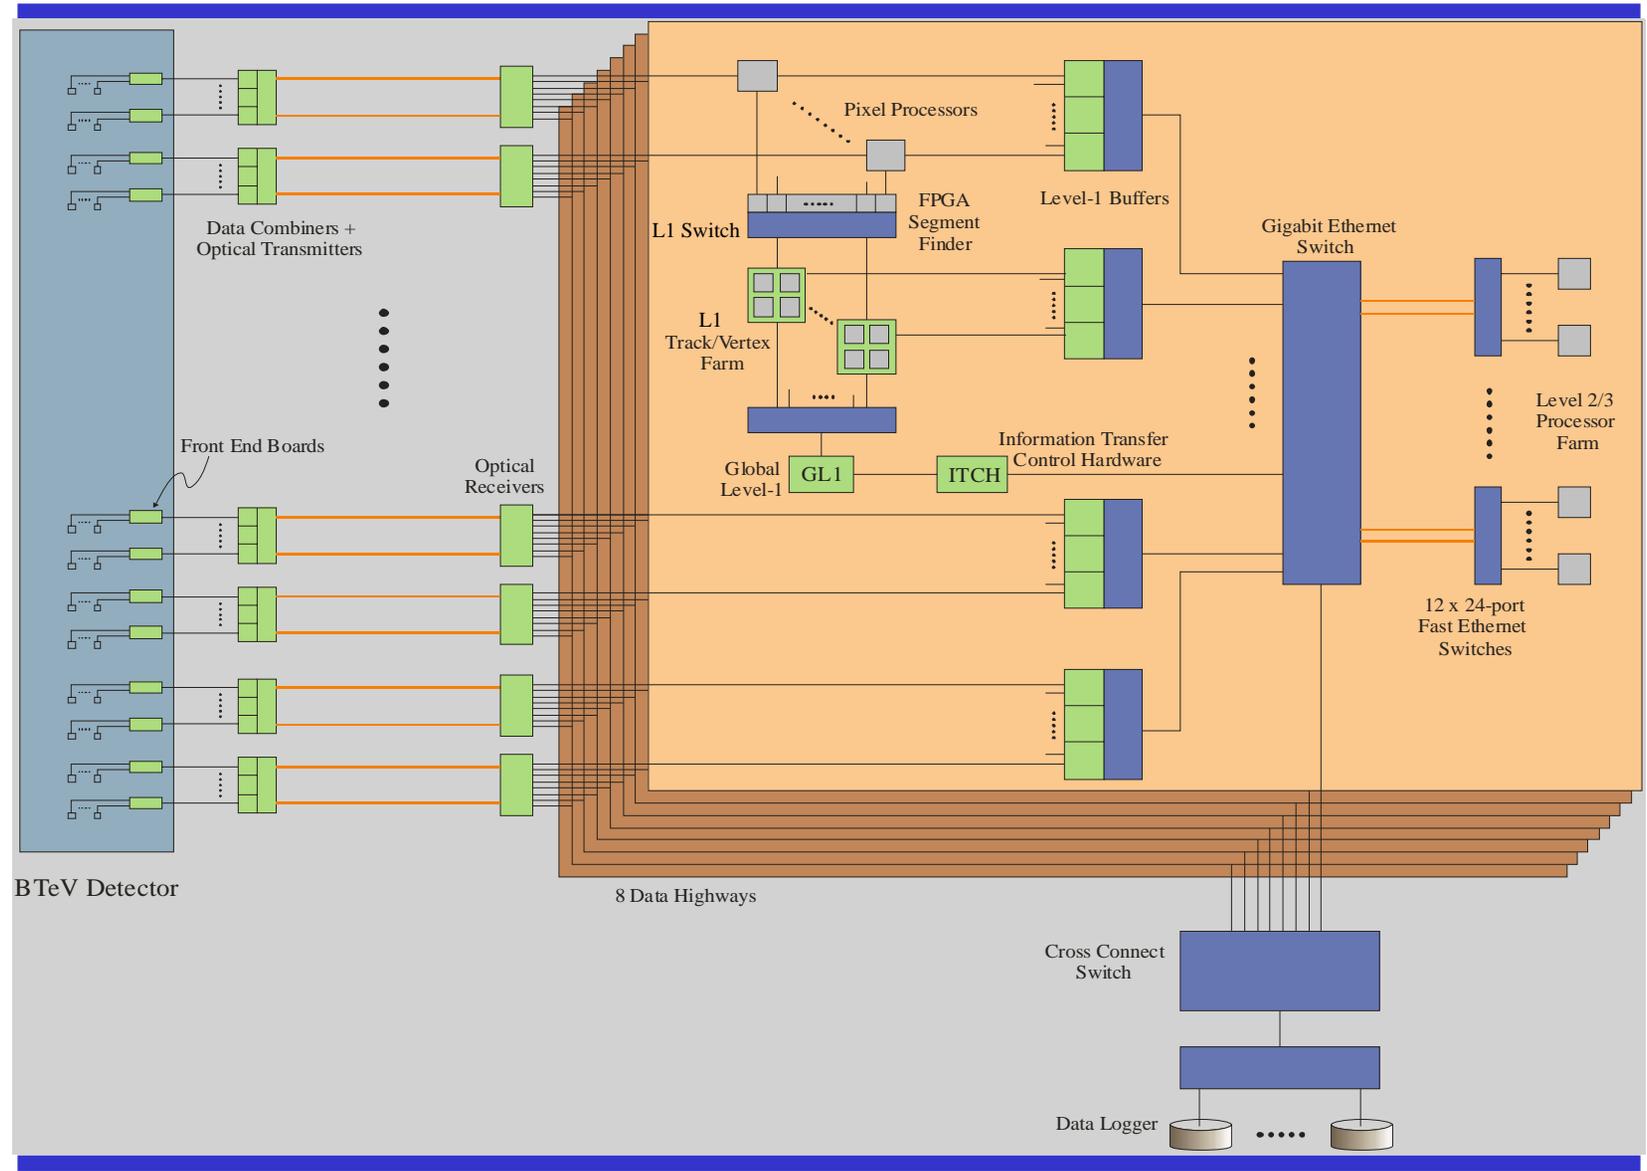
Base cost: \$11.1M (Material: \$7.9M, Labor: \$3.3M)

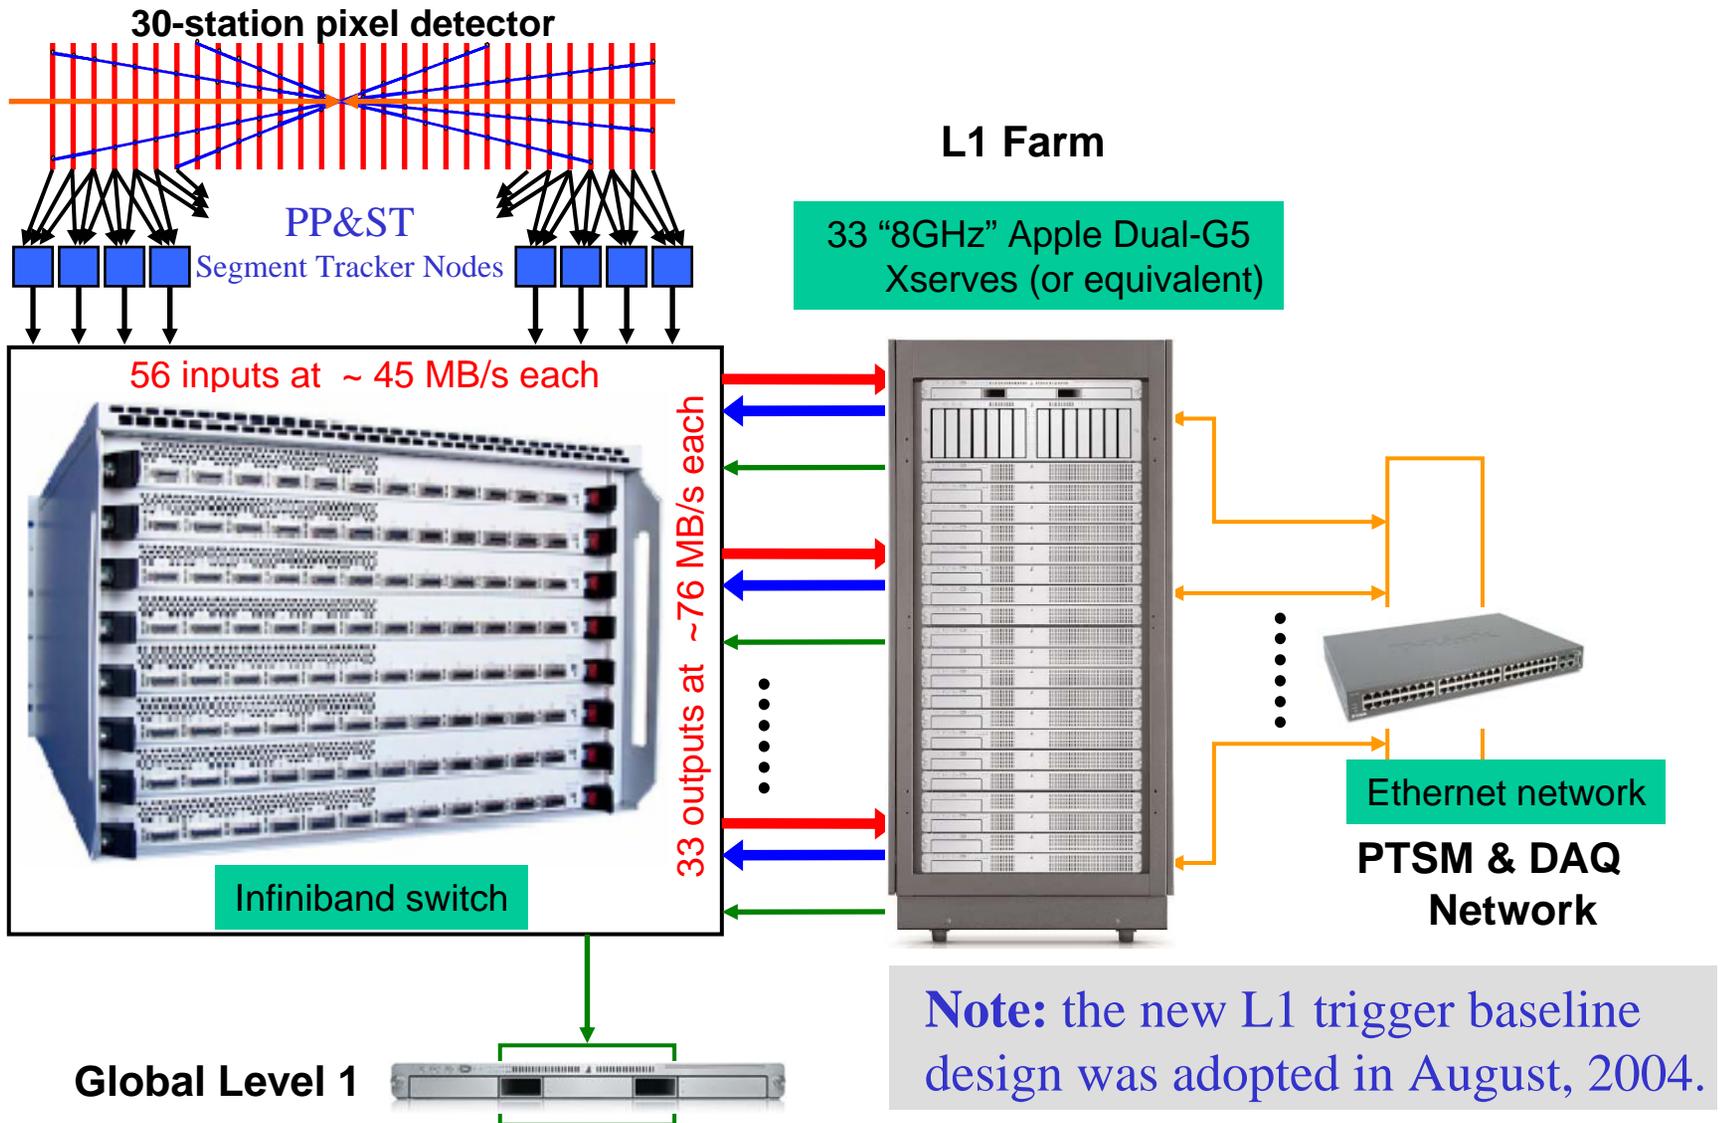


Base cost: \$11.1M (Material: \$7.9M, Labor: \$3.3M)



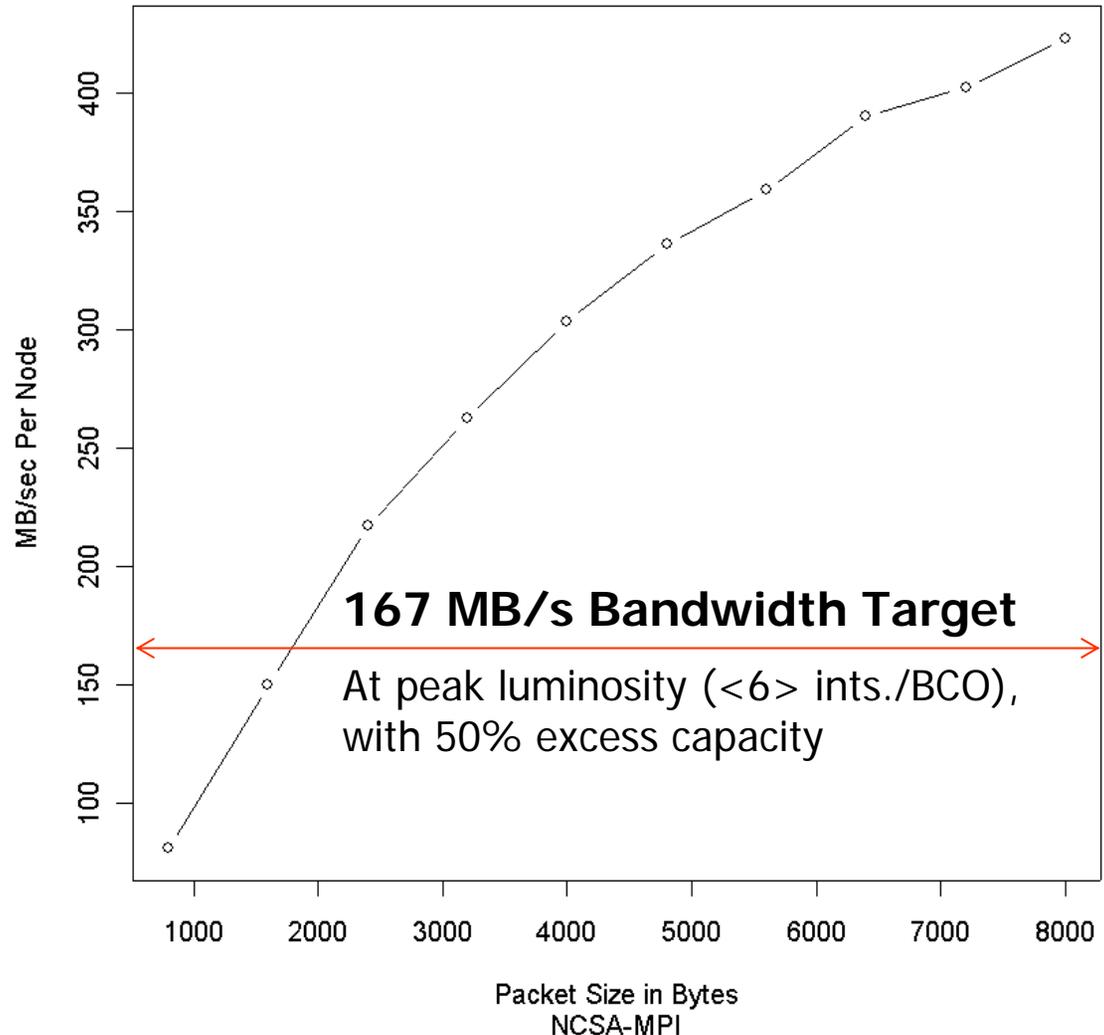
Three-level, eight highway trigger/DAQ architecture

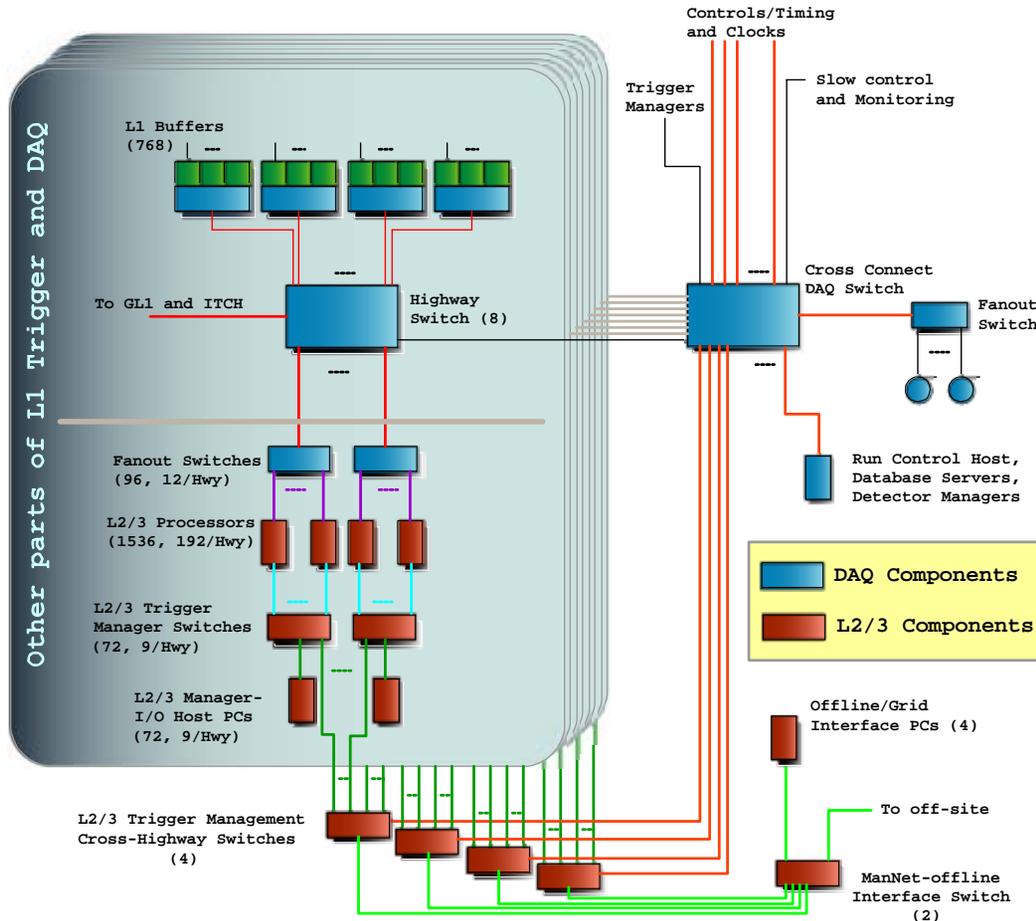




- The new baseline design for the L1 trigger includes an Infiniband switch, that replaces the custom switch we had in our previous baseline design.
- Bandwidth measurements confirm that an Infiniband switch **exceeds** bandwidth requirements for the L1 trigger.

Mean Transfer Speed per Segment Tracker Node in (17 -> 17)



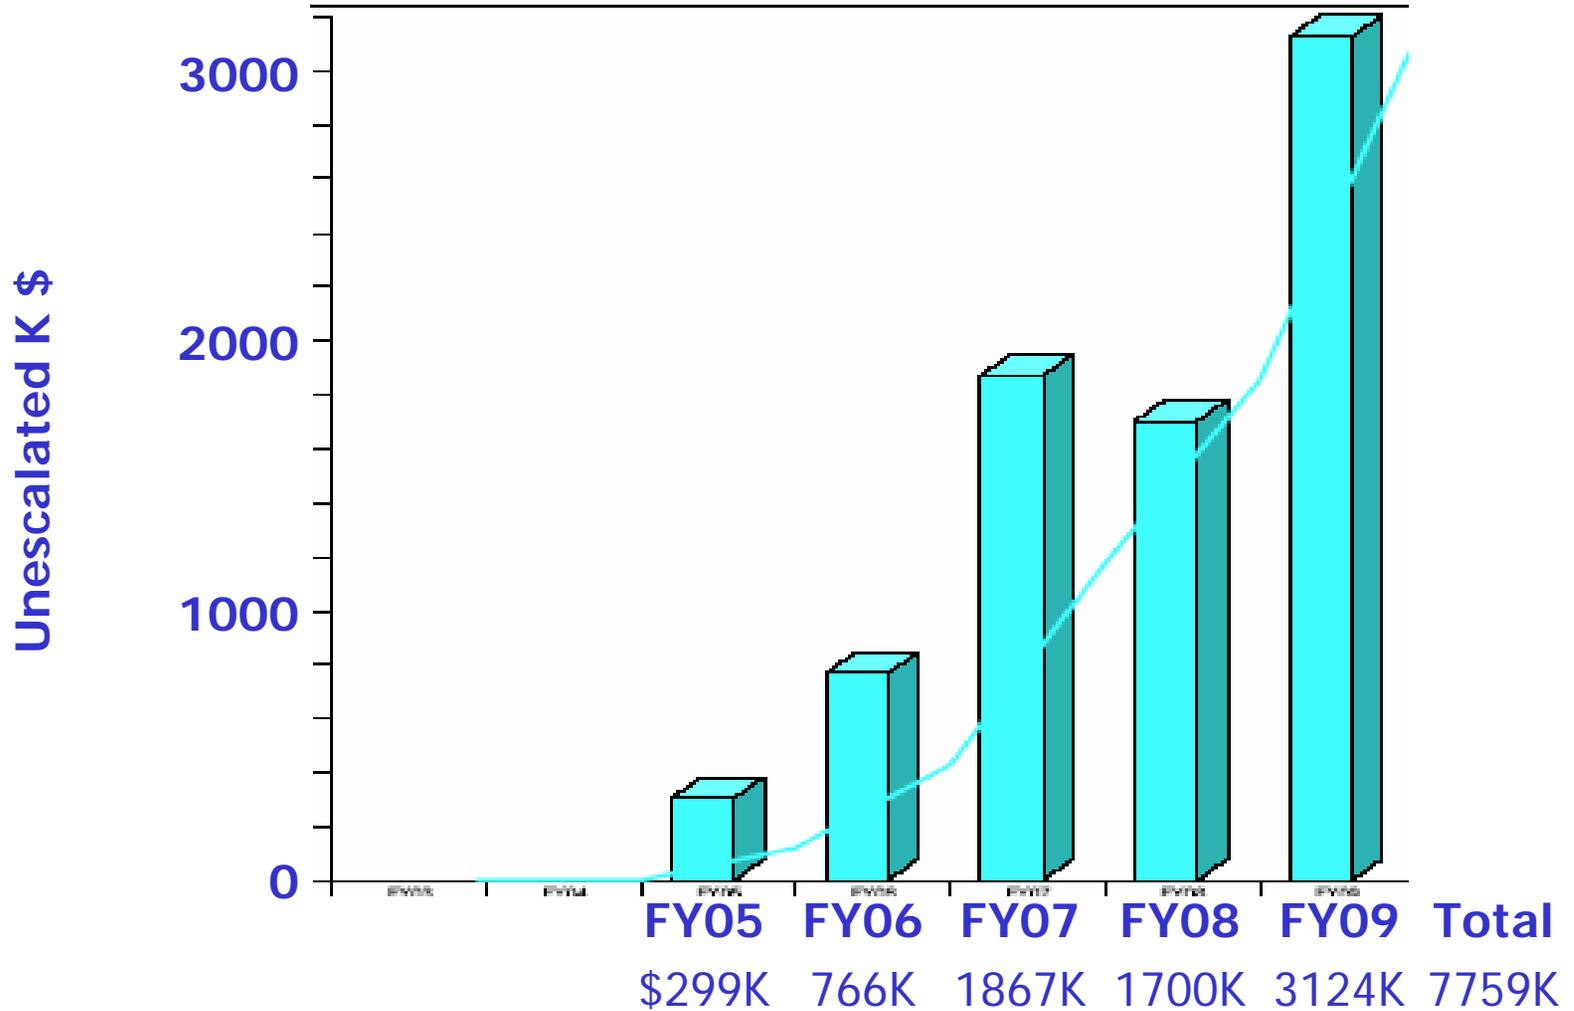


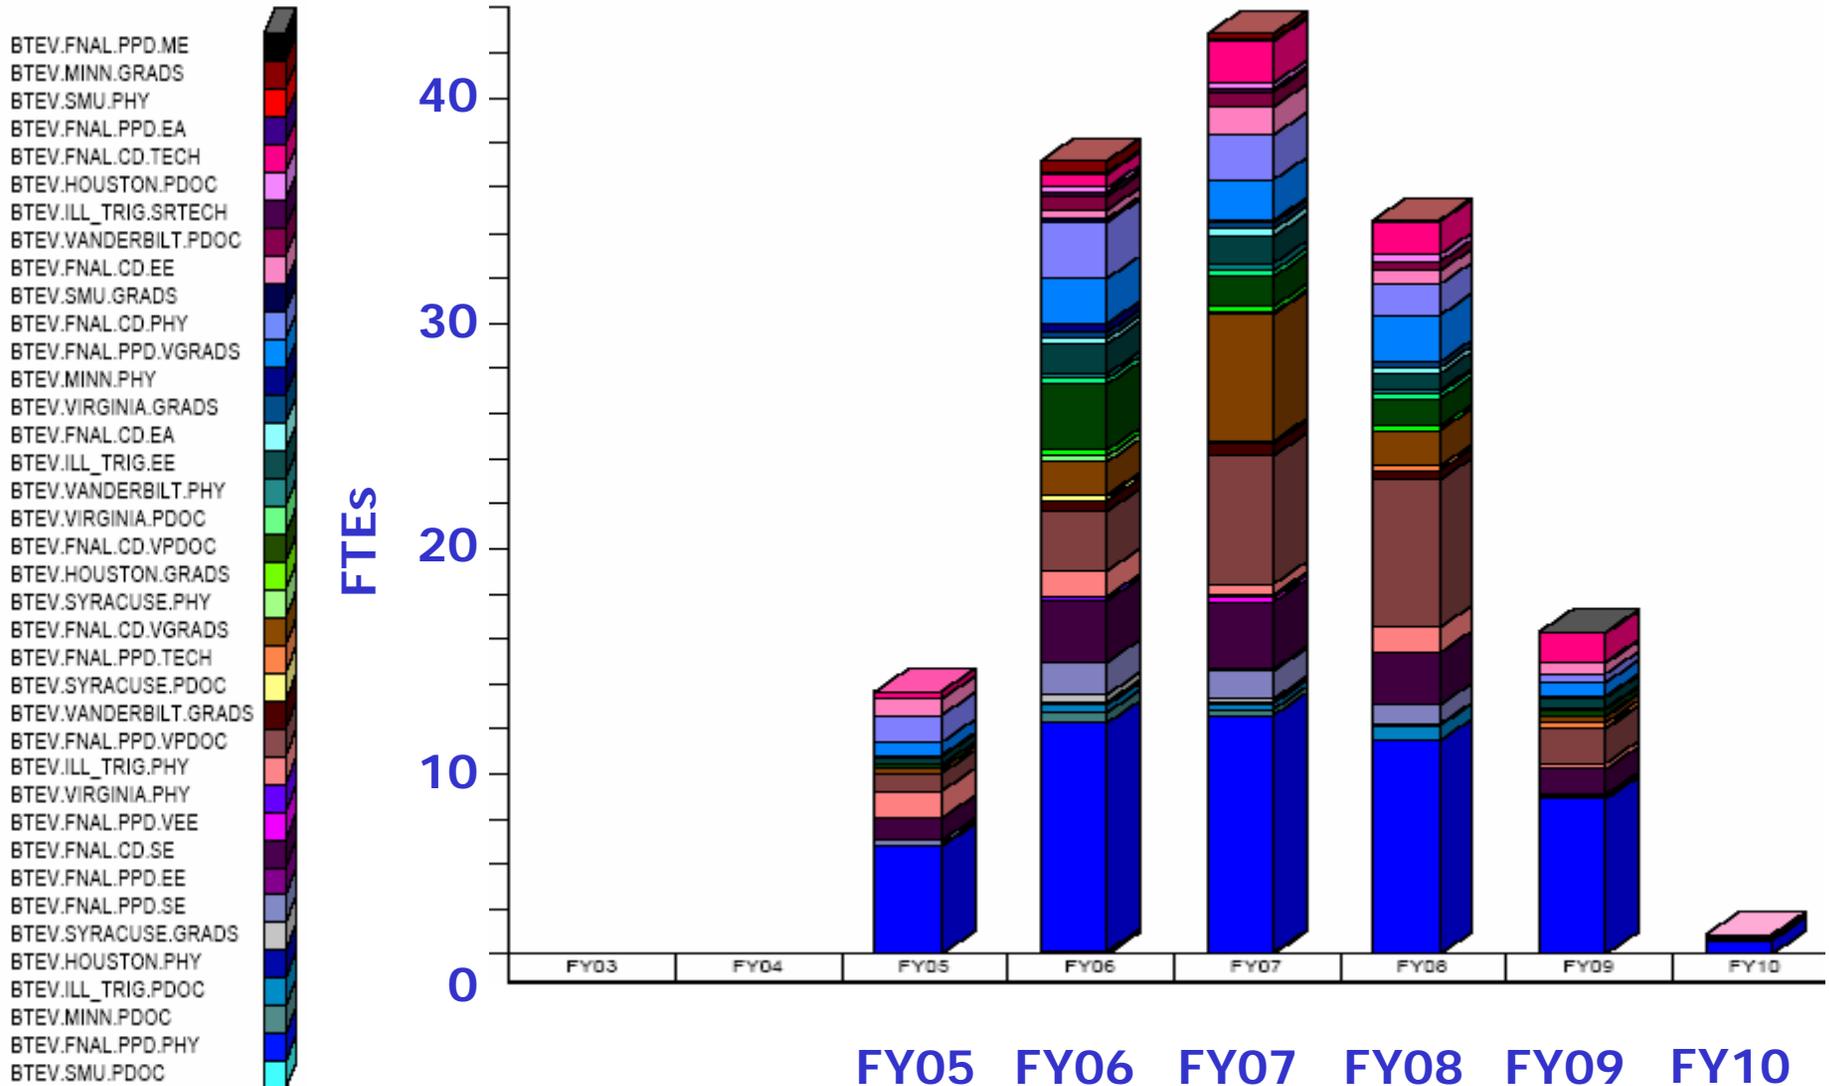
Baseline Design

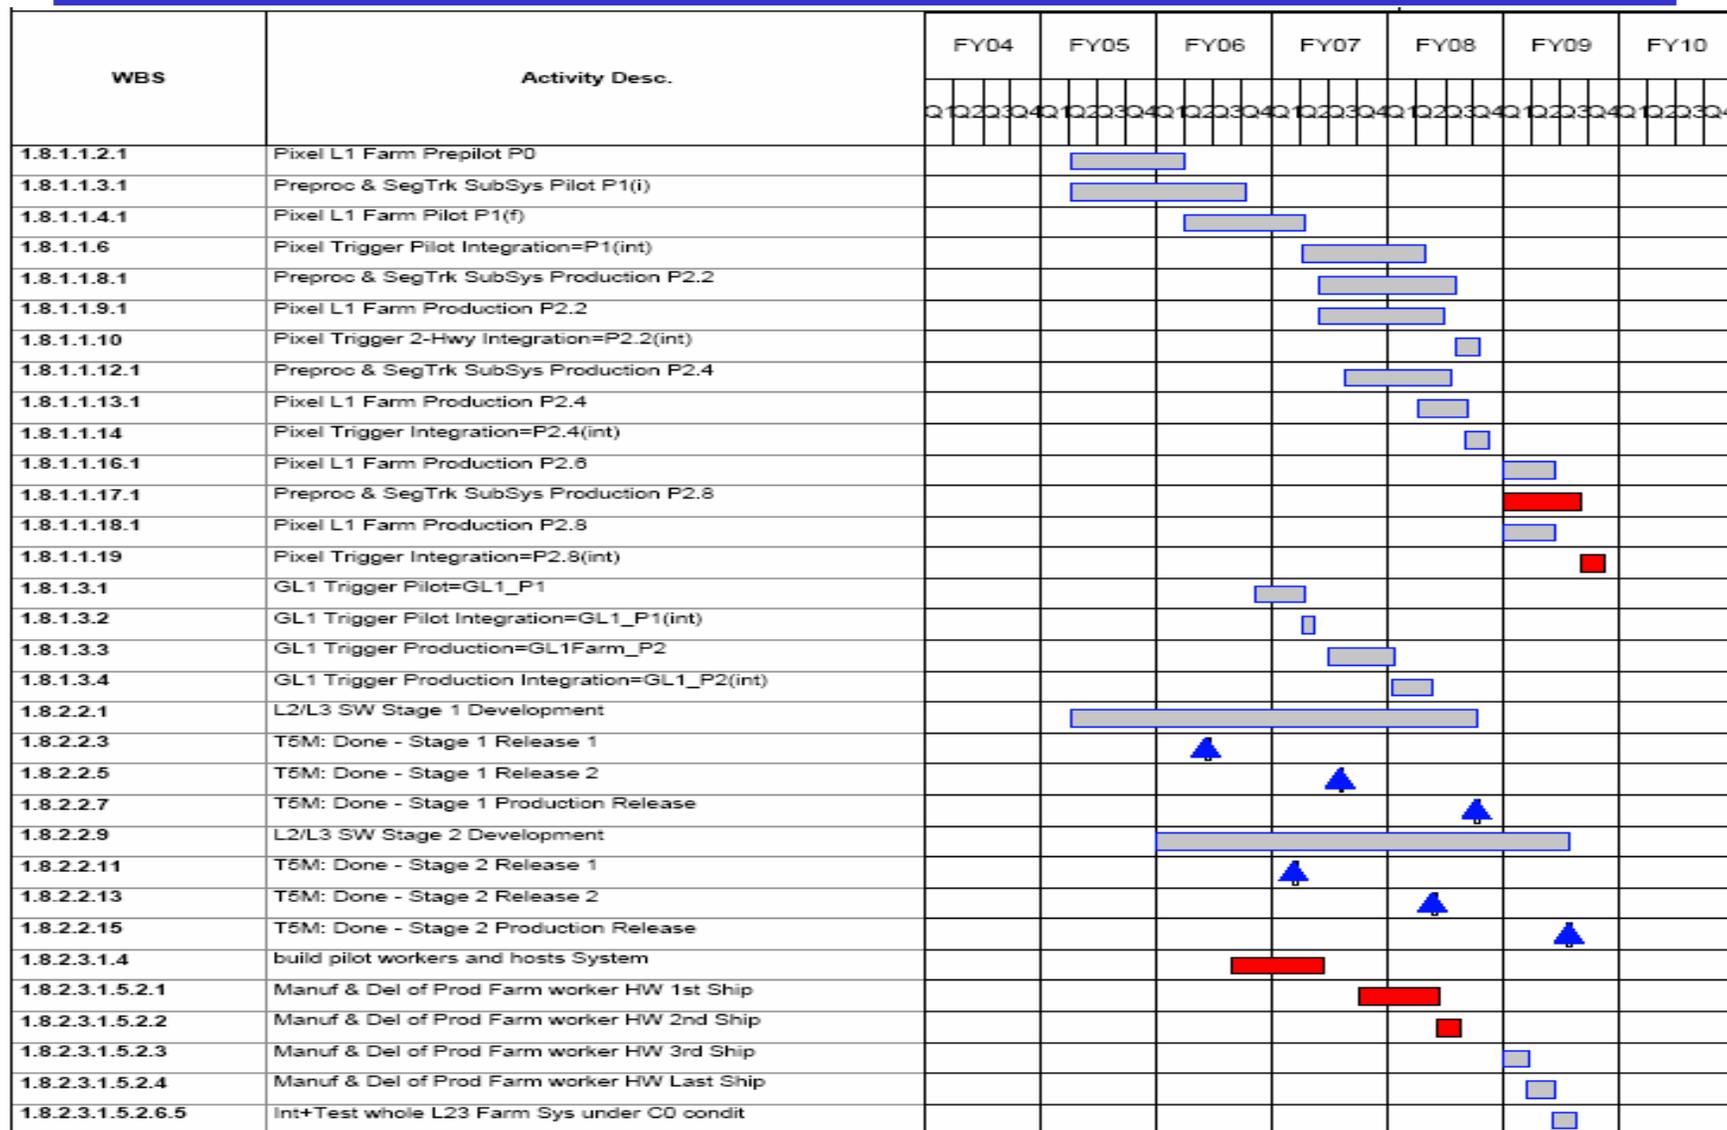
- L2/3 Processor farm consists of 1536 “12 GHz” CPUs (dual-CPU 1U rack-mount PCs)
- L2/3 trigger includes Manager-I/O Host PCs for database caches, worker management, monitoring, and event pool cache
- L2/3 Hardware in

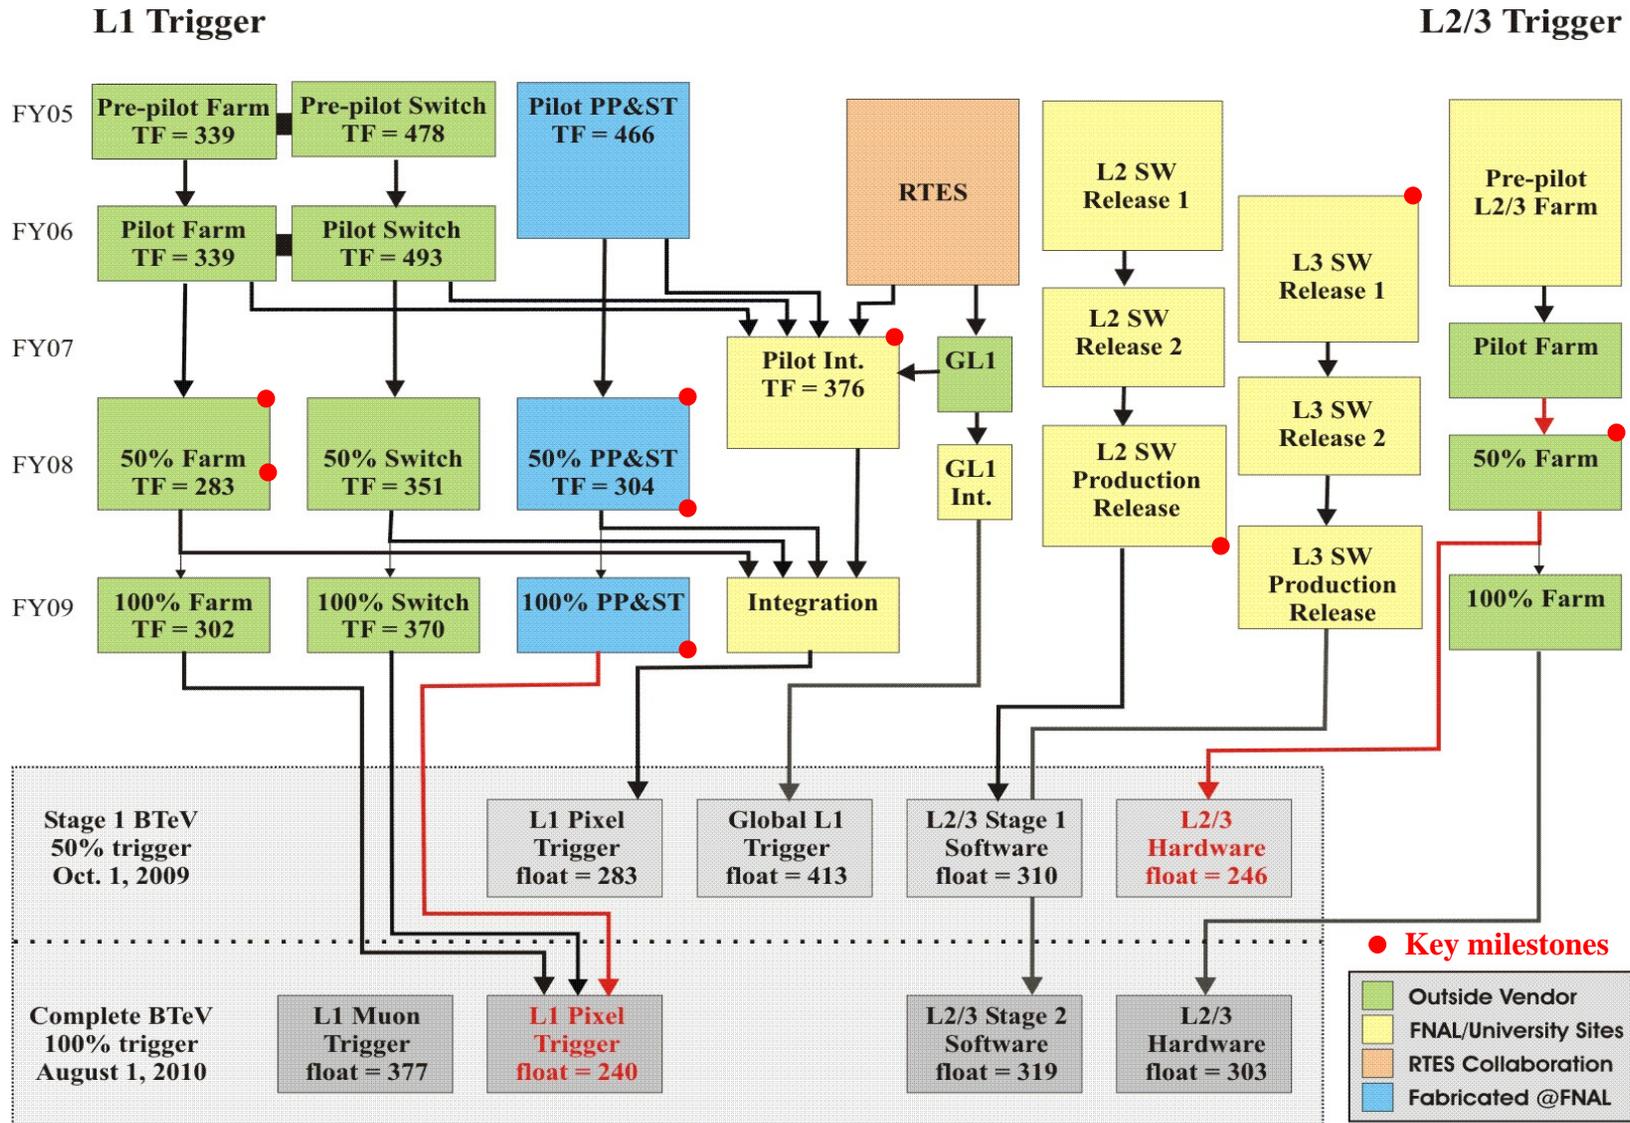
- L2 and L3 reconstruction software (tracks, vertices, photons, π^0 's, hyperons, neutral kaons, particle identification)
- L2 and L3 trigger algorithms
- Global L2 and Global L3 software (trigger lists & selection criteria)
- Alignment and calibration software
- Monitoring, feedback and event display software
- Software framework, utilities, and interfaces to databases
- DAQ interface software
- Offline filter and fast charm/beauty monitoring software (high-level filtering and monitoring software)

Activity ID	Activity Name	Base Cost (\$)	Material Contingency (%)	Labor Contingency (%)	Total FY05	Total FY06	Total FY07	Total FY08	Total FY09	Total FY10	Total FY05-10
1.8.1	L1 Hardware & Software	6,769,409	37	38	607,014	1,263,129	2,784,095	1,723,635	2,920,308	0	9,298,179
1.8.2	L2/L3 Hardware & Software	4,133,488	33	88	139,855	1,132,770	1,417,068	1,887,830	1,926,255	0	6,503,777
1.8.3	Trigger Electronics & SW Subproj Mgmt	230,199	16	24	46,879	54,336	53,697	53,697	53,484	19,178	281,271
1.8	file_18S_v10_091504	11,133,096	36	65	793,747	2,450,235	4,254,859	3,665,162	4,900,047	19,178	16,083,228









T2	Trigger pilot system tested	Jul-07
T2,T3	Stage 1 production release of L2/3 software	Mar-09
T3	Begin L1 2-highway pixel processor & segment tracker production	Jul-07
T3	End L1 2-highway pixel processor & segment tracker production	Dec-08
T3	Begin L1 2-highway farm production	Jul-07
T3	End L1 2-highway farm production	Oct-08
T3	Begin L2/3 farm worker node procurement	Jan-08
T3	Begin L3 software development	Oct-05
T3	Complete trigger system and integration with DAQ	Feb-10

- Stage 1 BTeV trigger (246 workdays of float)
 - Completing the first four highways for the L2/3 trigger has the lowest total float for the **Stage 1 BTeV detector**, with 246 days of float. The activities involve the procurement of computer farms. Procurement is delayed to obtain more processing capabilities for lower cost. The procurement of processors for the L2/3 trigger has minimal schedule risk, and there is considerable expertise for this work at Fermilab.
- Stage 2 BTeV trigger (240 workdays of float)
 - The completion of the remaining four highways of the L1 pixel preprocessor and segment tracker (PP&ST) hardware has the lowest total float for the **Stage 2 BTeV detector**, with 240 days of float. By the time this work is started we will have considerable expertise building, testing, and commissioning PP&ST hardware. Therefore, we expect minimal schedule risk, since four trigger highways will be fully operational by the time this work begins.

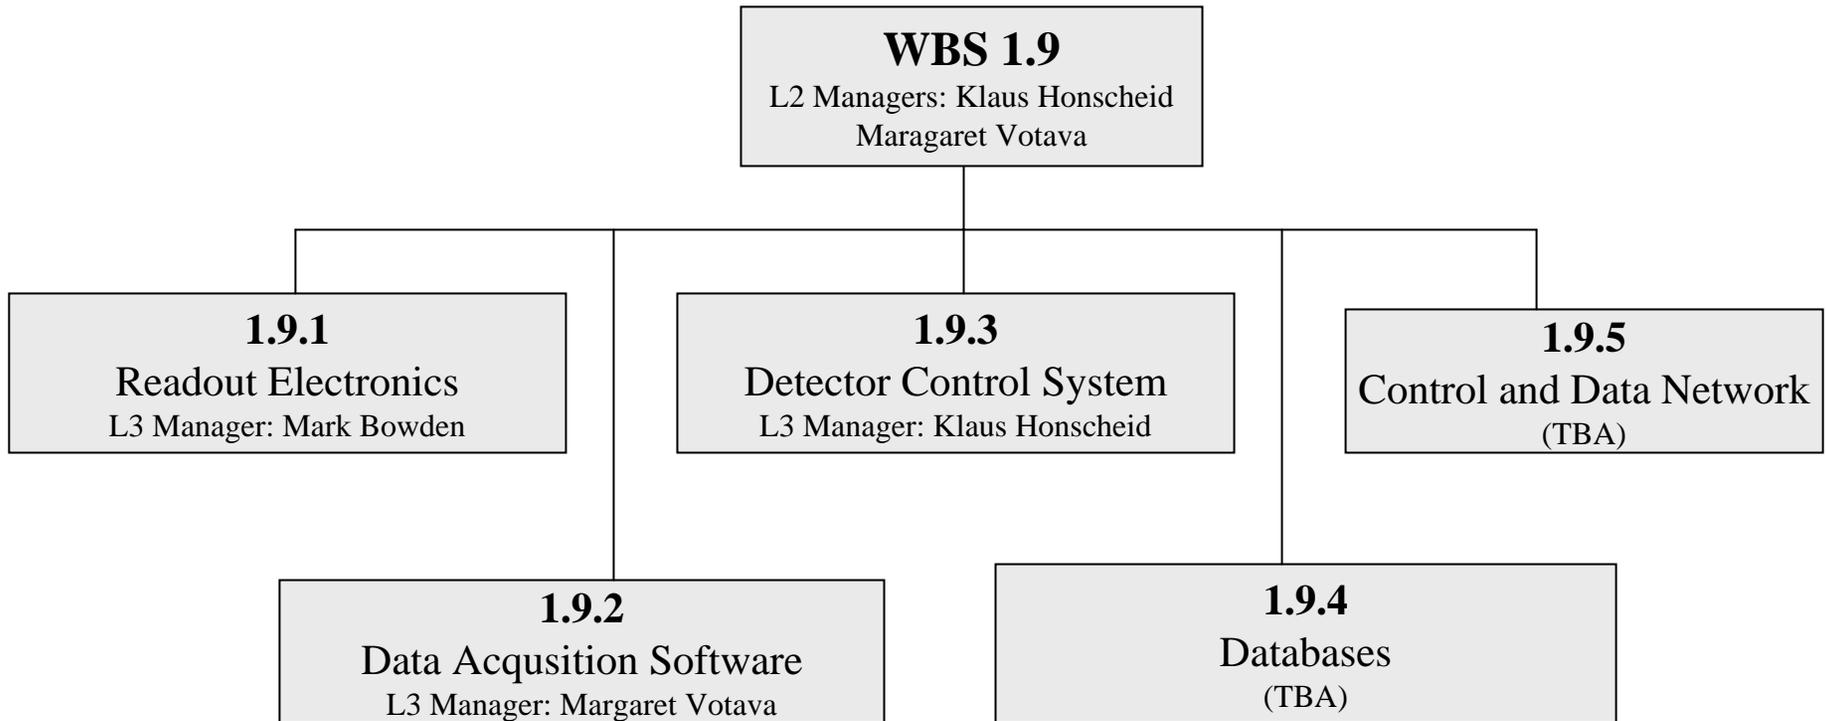
WBS Item	Risk Event	Response/Mitigation Strategy
For example: 1.8.1.1.2.1, 1.8.1.1.4.1	Baseline processor fails to meet requirements	Benchmark and qualify more than one processor during R&D or early construction phase. Have 2 nd option ready if 1 st option is unsatisfactory.
For example: 1.8.1.1.8.1.1, 1.8.1.2.1.2	L1 pixel segment tracker or muon preprocessor algorithm exceeds the size of the selected FPGA. Larger FPGA increases the cost	Use contingency funds to upgrade to a bigger and more expensive FPGA. Consider other implementation alternatives early in the design stage. Consider simplifying the algorithms.
For example: 1.8.1.1.3.1.1, 1.8.1.1.8.1.1	Communication links do not satisfy error rate specifications	Use contingency to modify PC boards, buy new parts, connectors, or cables.
For example: 1.8.1.1.4.1.1.5, 1.8.2.2	Shortage of software developers	Prioritize critical tasks. Use contingency to hire software developers temporarily.
Numerous WBS items	Experienced people leave	Be sure that more than one person is working on critical tasks. Use contingency funds to hire the person leaving as a temporary consultant while their expertise is transferred to existing or new personnel.

- Develop a schedule which (a) completes critical design and validation activities as soon as possible and is ready for production six to nine months in advance of the production start date, and (b) completes production of the trigger and data acquisition systems six to nine months in advance of first collisions.
 - (a) Critical design and validation activities have been an ongoing effort. We will complete the L1 PP&ST system 8 months before the start of production.
 - (b) We have developed a schedule that completes 50% of the L1 trigger more than 13 months before the need-by date for the Stage 1 detector, and completes 50% of the L2/3 trigger almost one year before the need-by date.
- Re-evaluate the basis of estimate of the FPGA costs to allow for uncertainty in the de-escalation profile.
 - We no longer de-escalate FPGA costs.
- Quickly identify and apply new individuals and groups to provide the physicist effort for by the WBS.
 - We have identified new individuals and groups (Univ. of Houston, Southern Methodist University, Univ. of Virginia), and will continue to do so.

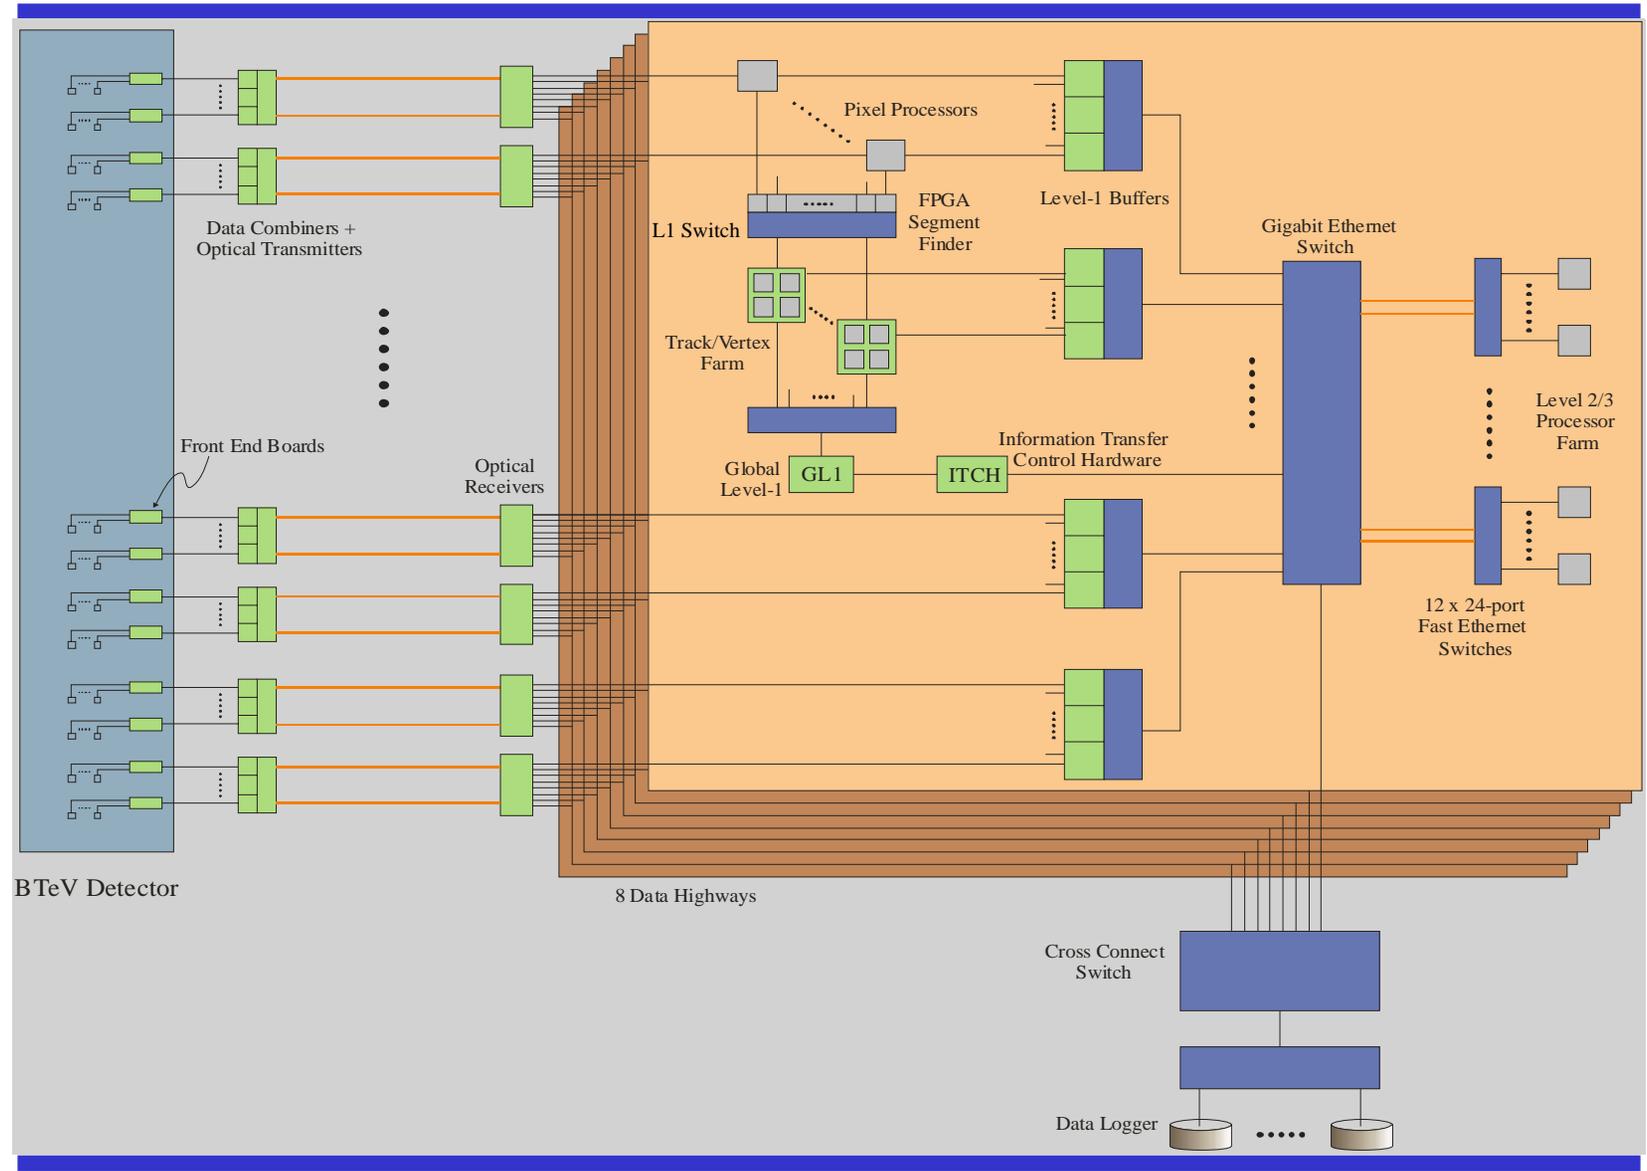
- Readout Electronics
- Data Acquisition Software
- Detector Control System
- Databases
- Control & Data Network

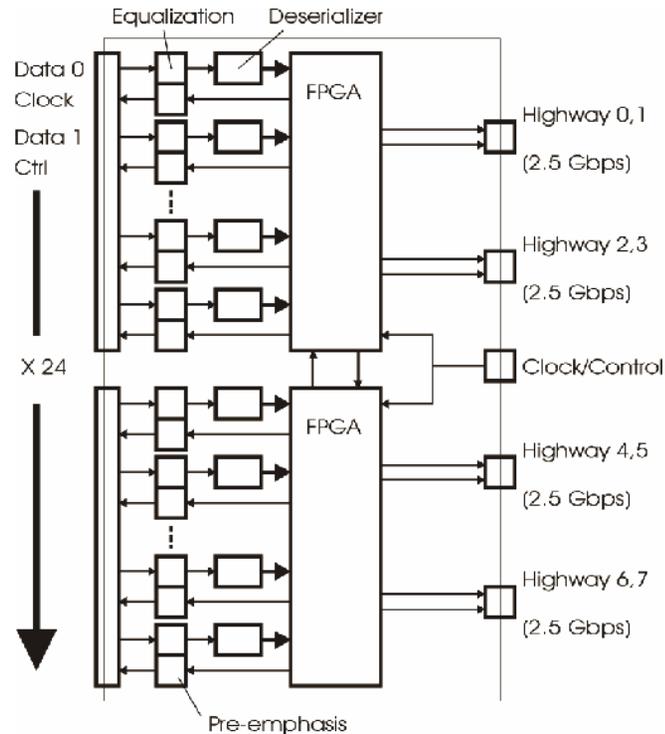
Base cost: \$12.7M (Material: \$6.0M, Labor: \$6.7M)

Base cost: \$12.7M (Material: \$6.0M, Labor: \$6.7M)

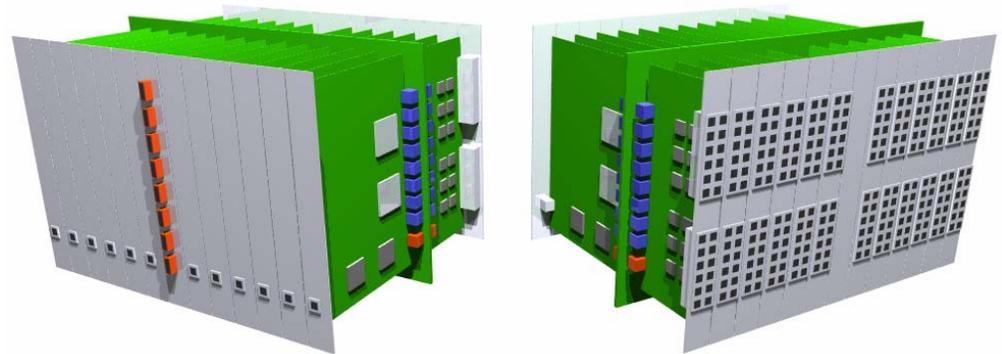


Three-level, eight highway trigger/DAQ architecture

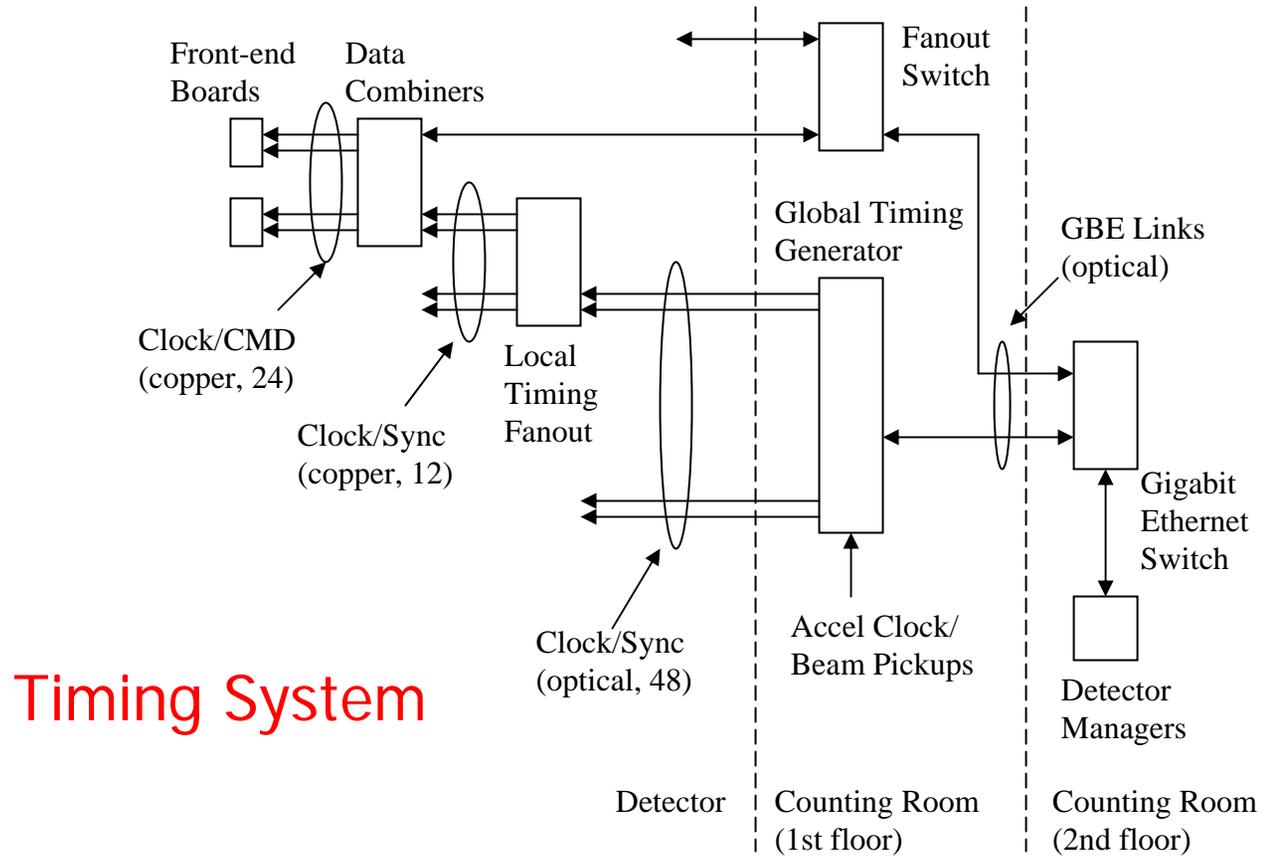




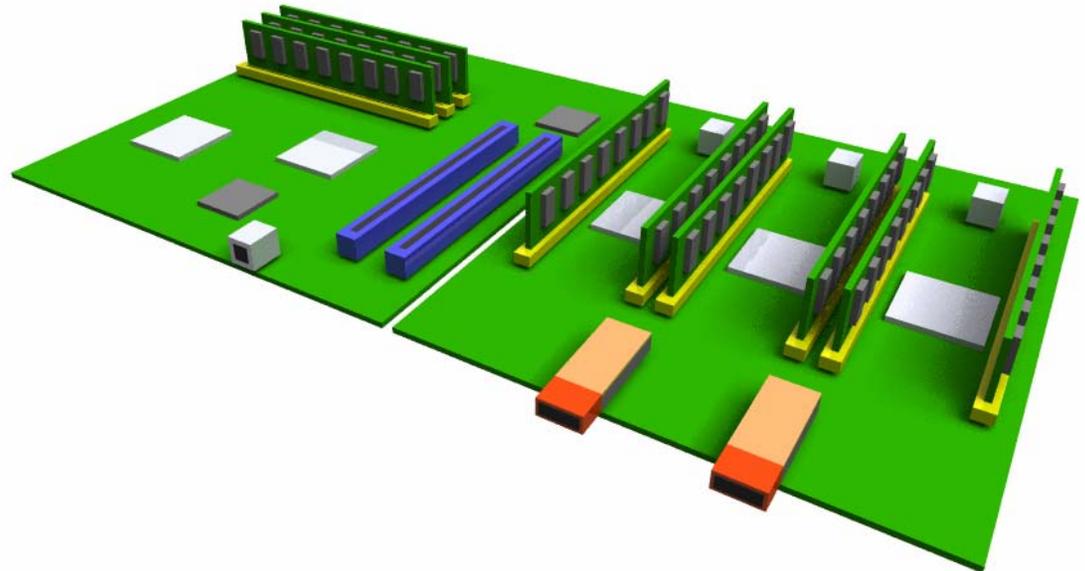
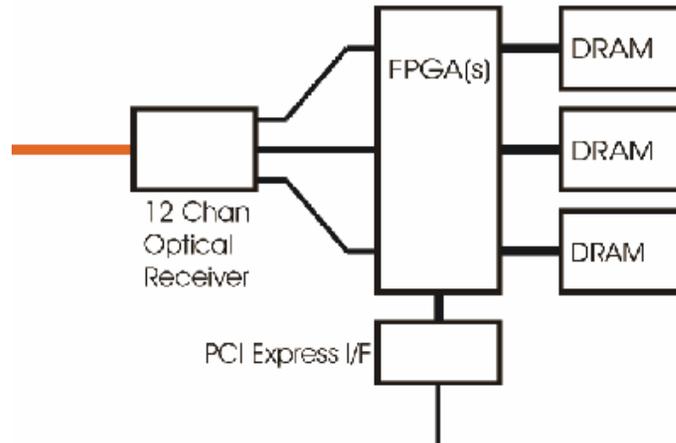
Data Combiner



Input receiver/multiplexer for detector front-end boards.



Fast control and timing network for precise system synchronization.

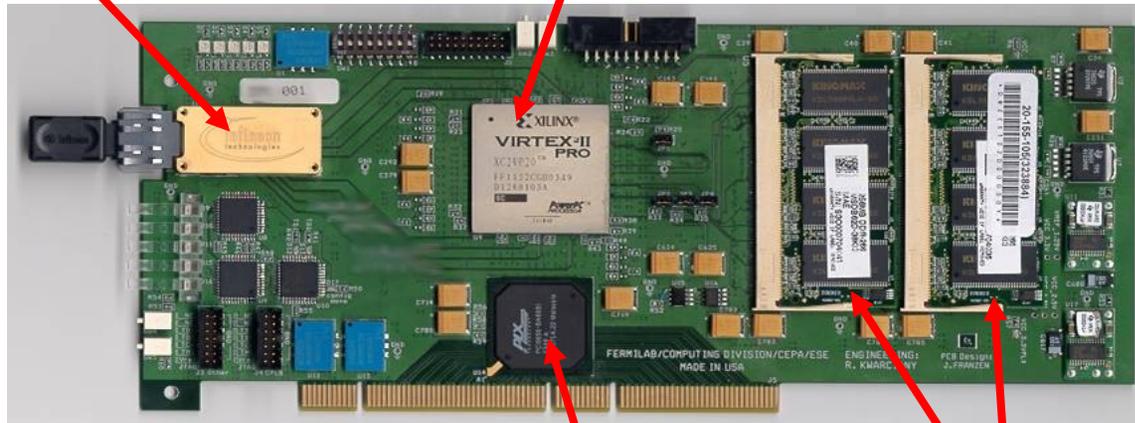


Large capacity buffers to hold detector data while L1 is processing pixel & muon data.

Optical Receiver

(8 channels @ 2.5 Gbps)

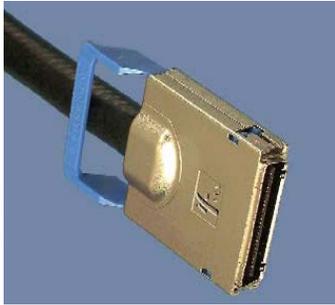
FPGA (deserializers, memory controller)



DRAM (256MB X 2,
DDR SODIMM)

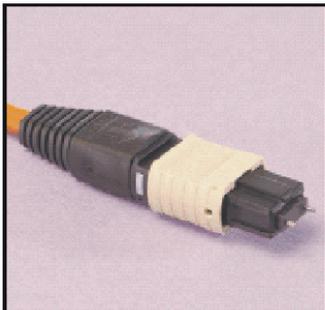
PCI Interface

- Standard optical interface
- Large buffer memory
- PCI Interface (BTeV will use a newer generation, e.g. based on PCI-Express)



Front-end to Data Combiners
- 600 Mbps copper

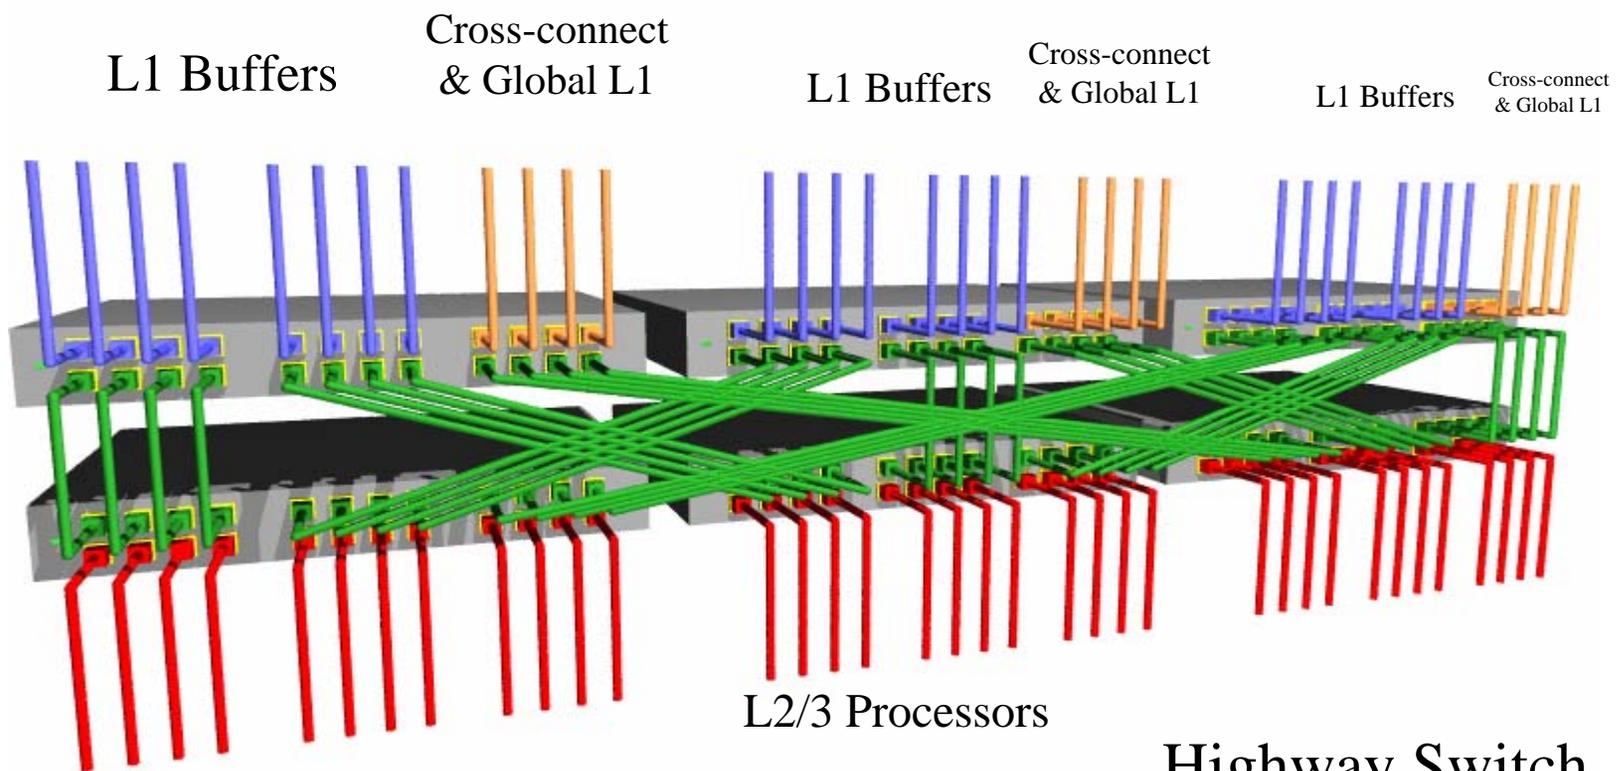
**Data Combiners to L1 buffers
& L1 Trigger**
- 2.5 Gbps optical



L1 Trigger to L1 Buffers
- 2.5 Gbps copper

Network
- 1 Gbps (CAT6) copper

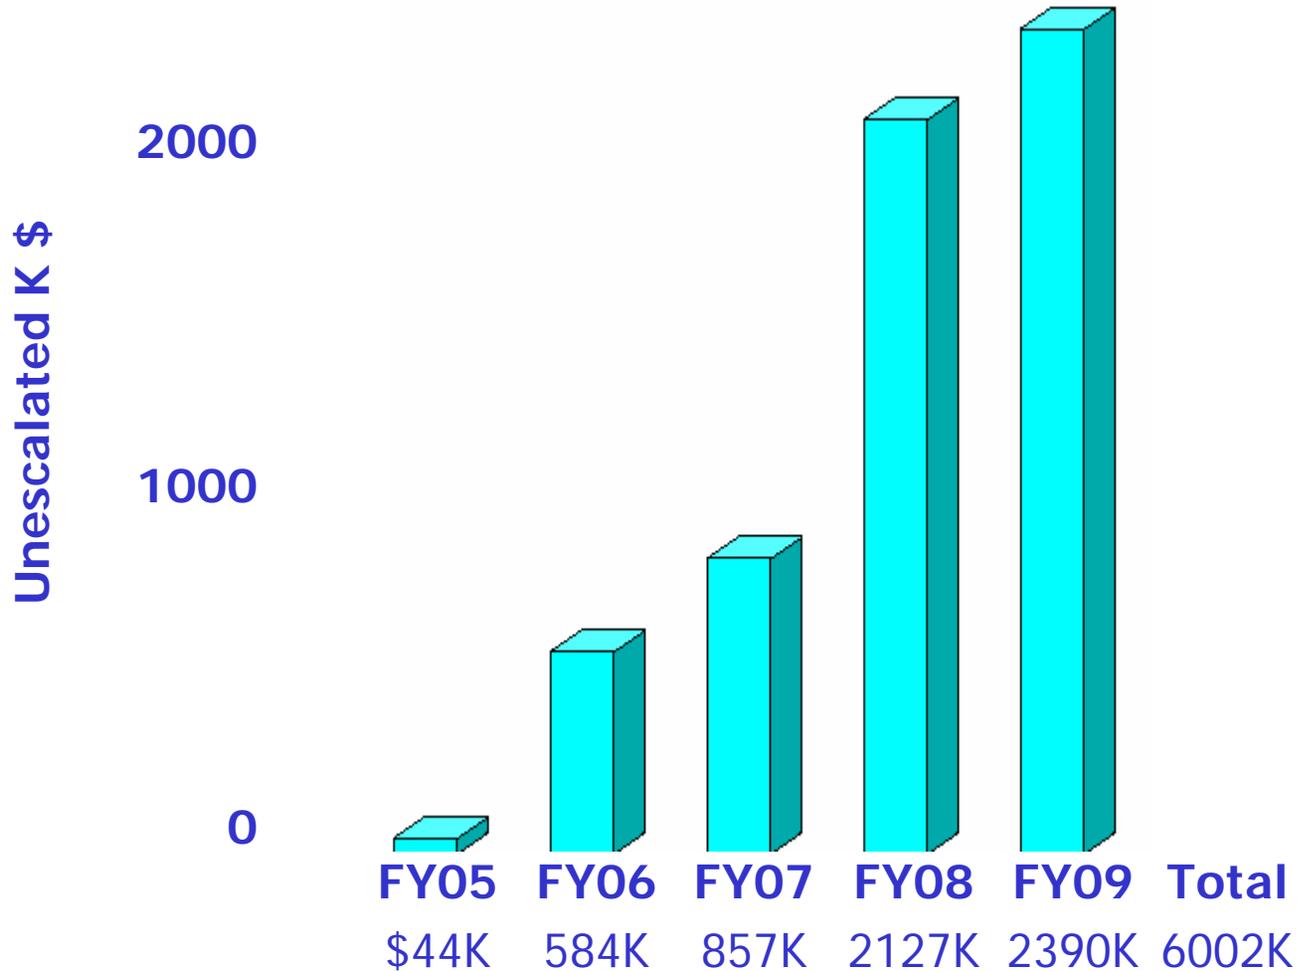
(all point-to-point serial)

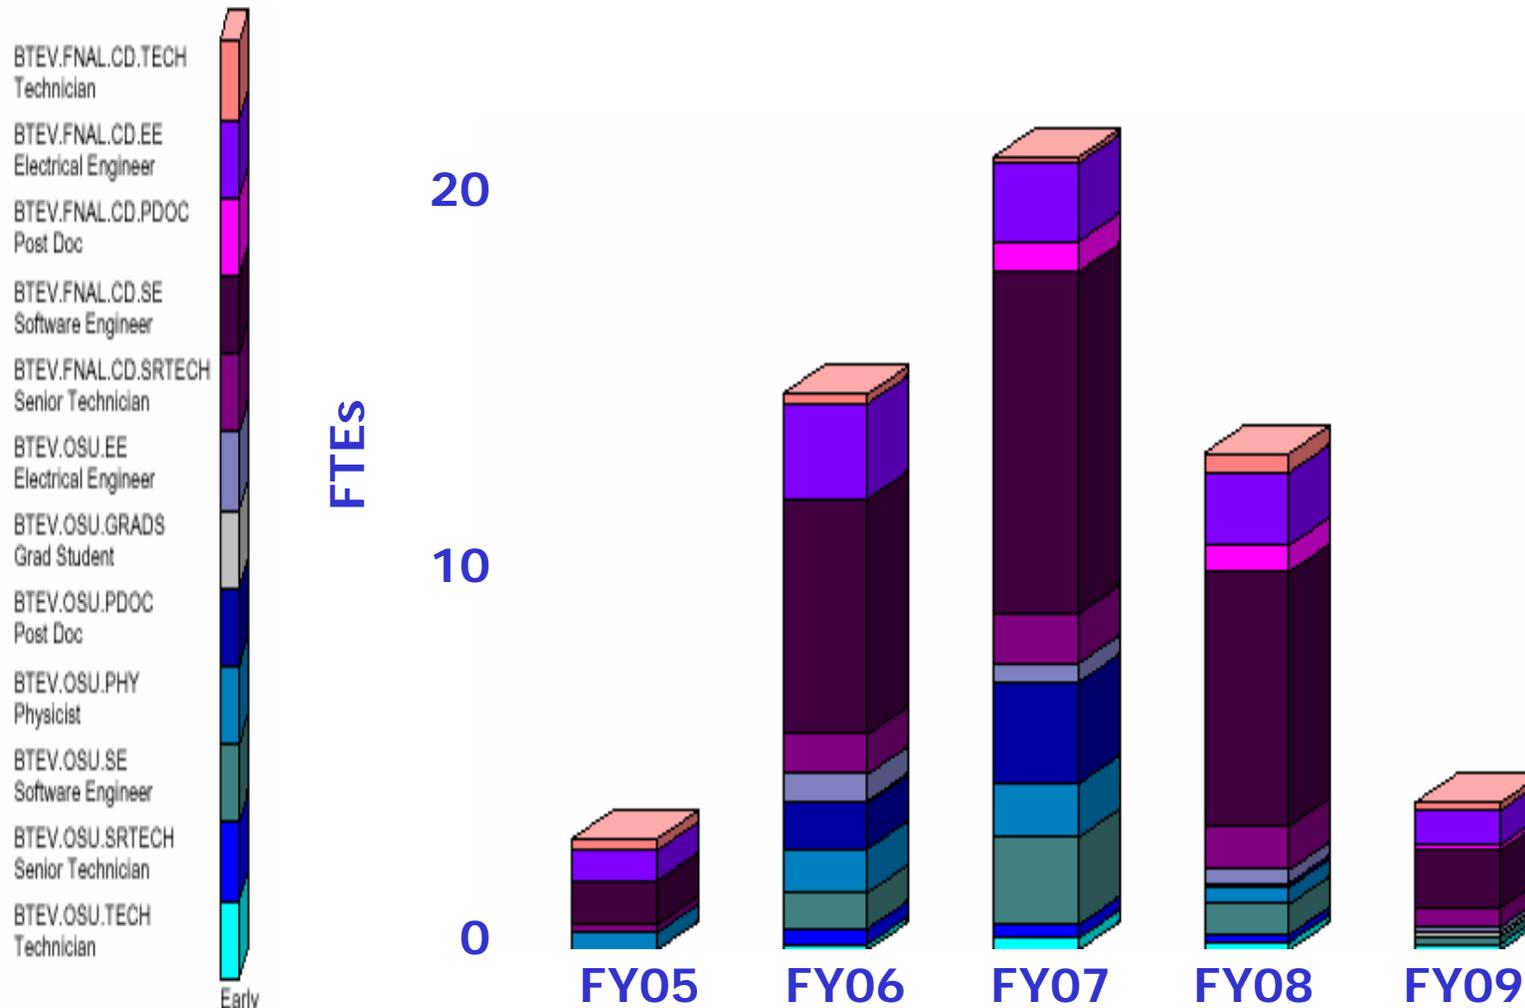


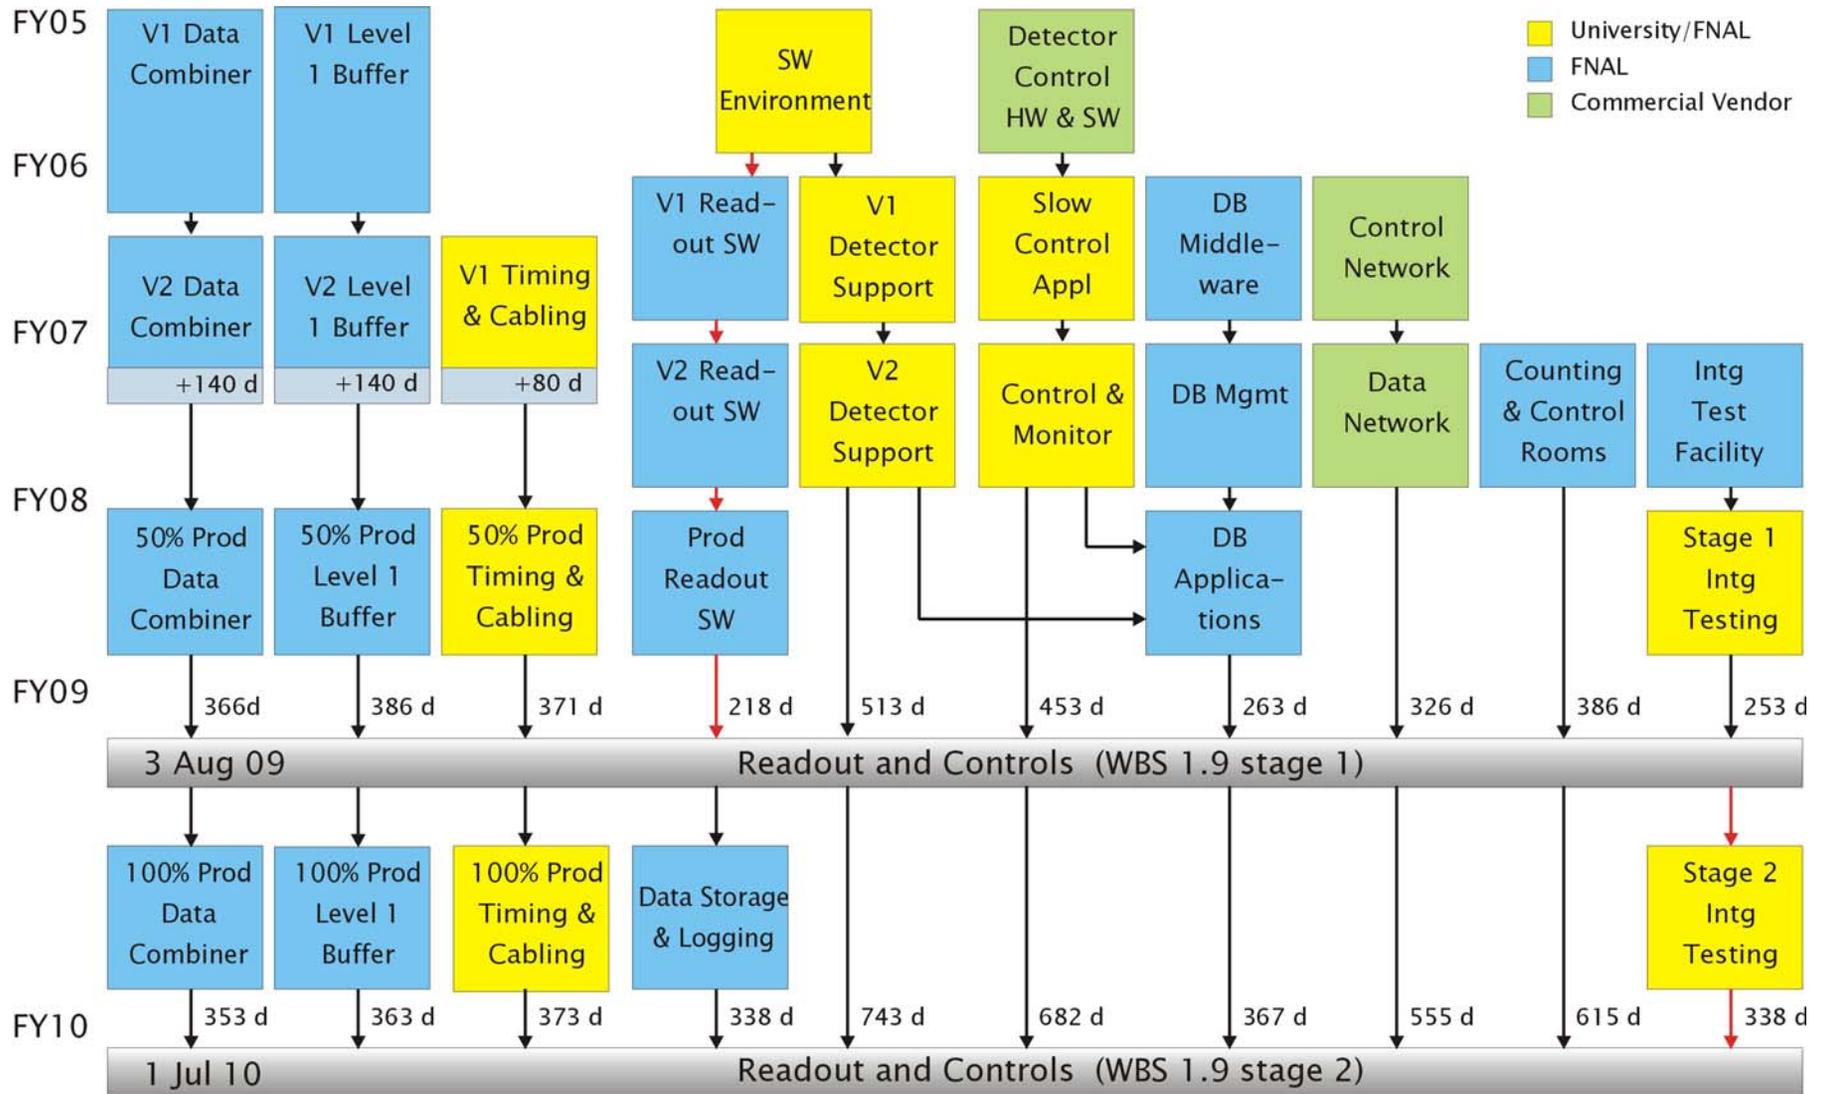
Highway Switch
72 Port Gigabit Ethernet

- Configuration Subsystem - software to download, initialize and partition all system components
- Run Control Subsystem - software to control and monitor the operation and overall dataflow of the system
- Detector Control - “slow” control network to set and monitor all system environmental parameters
- Databases - store and access operating parameters, maintain a time history of all system variables, and store and access parameters necessary for trigger algorithms
- Infrastructure - counting and control room infrastructure

Activity ID	Activity Name	Base Cost (\$)	Material Contingency (%)	Labor Contingency (%)	Total FY05	Total FY06	Total FY07	Total FY08	Total FY09	Total FY10	Total FY05-10
1.9.1	Readout Electronics	5,576,109	43	30	285,477	889,953	1,273,413	2,720,348	2,585,086	0	7,754,277
1.9.2	Data Acquisition Software	2,326,663	39	30	203,082	592,416	850,937	361,018	1,093,241	0	3,100,694
1.9.3	Detector Control System	396,595	24	30	0	197,735	301,057	4,950	0	0	503,742
1.9.4	Databases	1,473,399	60	29	0	535,591	800,334	592,636	23,262	0	1,951,823
1.9.5	Control & Data Network	334,986	60	30	0	148,027	60,622	271,655	0	0	480,304
1.9.6	Infrastructure & Integration	973,725	31	30	0	34,384	434,070	760,413	40,282	0	1,269,149
1.9.7	Technical Support Activities	1,029,081	31	28	25,829	415,500	219,552	279,506	278,486	109,104	1,327,976
1.9.8	Readout & Controls Subproject Management	604,858	30	21	11,809	214,337	235,447	234,977	36,847	0	733,417
1.9	file_19s_092004	12,715,417	41	29	526,197	3,027,942	4,175,431	5,225,501	4,057,204	109,104	17,121,380







T2, T3	Data Combiner Board pre-production units tested and approved	Aug-06
T2	Multinode release of Data Acquisition RCS package	Mar-08
T3	Production DCB delivered and tested	Feb-09
T3	Production Level 1 Buffers delivered and tested	Jan-09
T3	Single node release of RCS package	Sep-07
T3	Data Acquisition software complete	Jun-09
T3	Calibration and Trigger database complete	Oct-08

- Stage II has more than 330 days of total float, so the critical path for WBS 1.9 is in the Stage I deliverables:
 - Partitioning release of run control (total float = 218 workdays)
 - This is the final release of run-control software that includes the ability to partition portions of the detector to acquire data independently for BTeV commissioning.
 - Database Application completion (total float = 263 workdays)
 - Completion of the database applications that will be needed for data taking. This includes running conditions, trigger databases, and calibration databases.

- Develop a schedule which completes critical design and validation activities as soon as possible and is ready for production 6 to 9 months in advance of the production start date.
 - Pixel DCB production has been moved from WBS 1.2 to WBS 1.9. This allows us to begin the design effort for all DCBs in FY05, which is one year earlier than in our previous schedule. We will complete the DCB design 7 months before the start of production.

- Re-evaluate the bases of estimate of the FPGA costs to allow for uncertainty in the de-escalation profile.
 - We no longer de-escalate electronics costs. For some components, we assume a nominal increase in performance between now and the time of purchase.

More information on the trigger (WBS 1.8) and DAQ (WBS 1.9) is available in the breakout sessions.

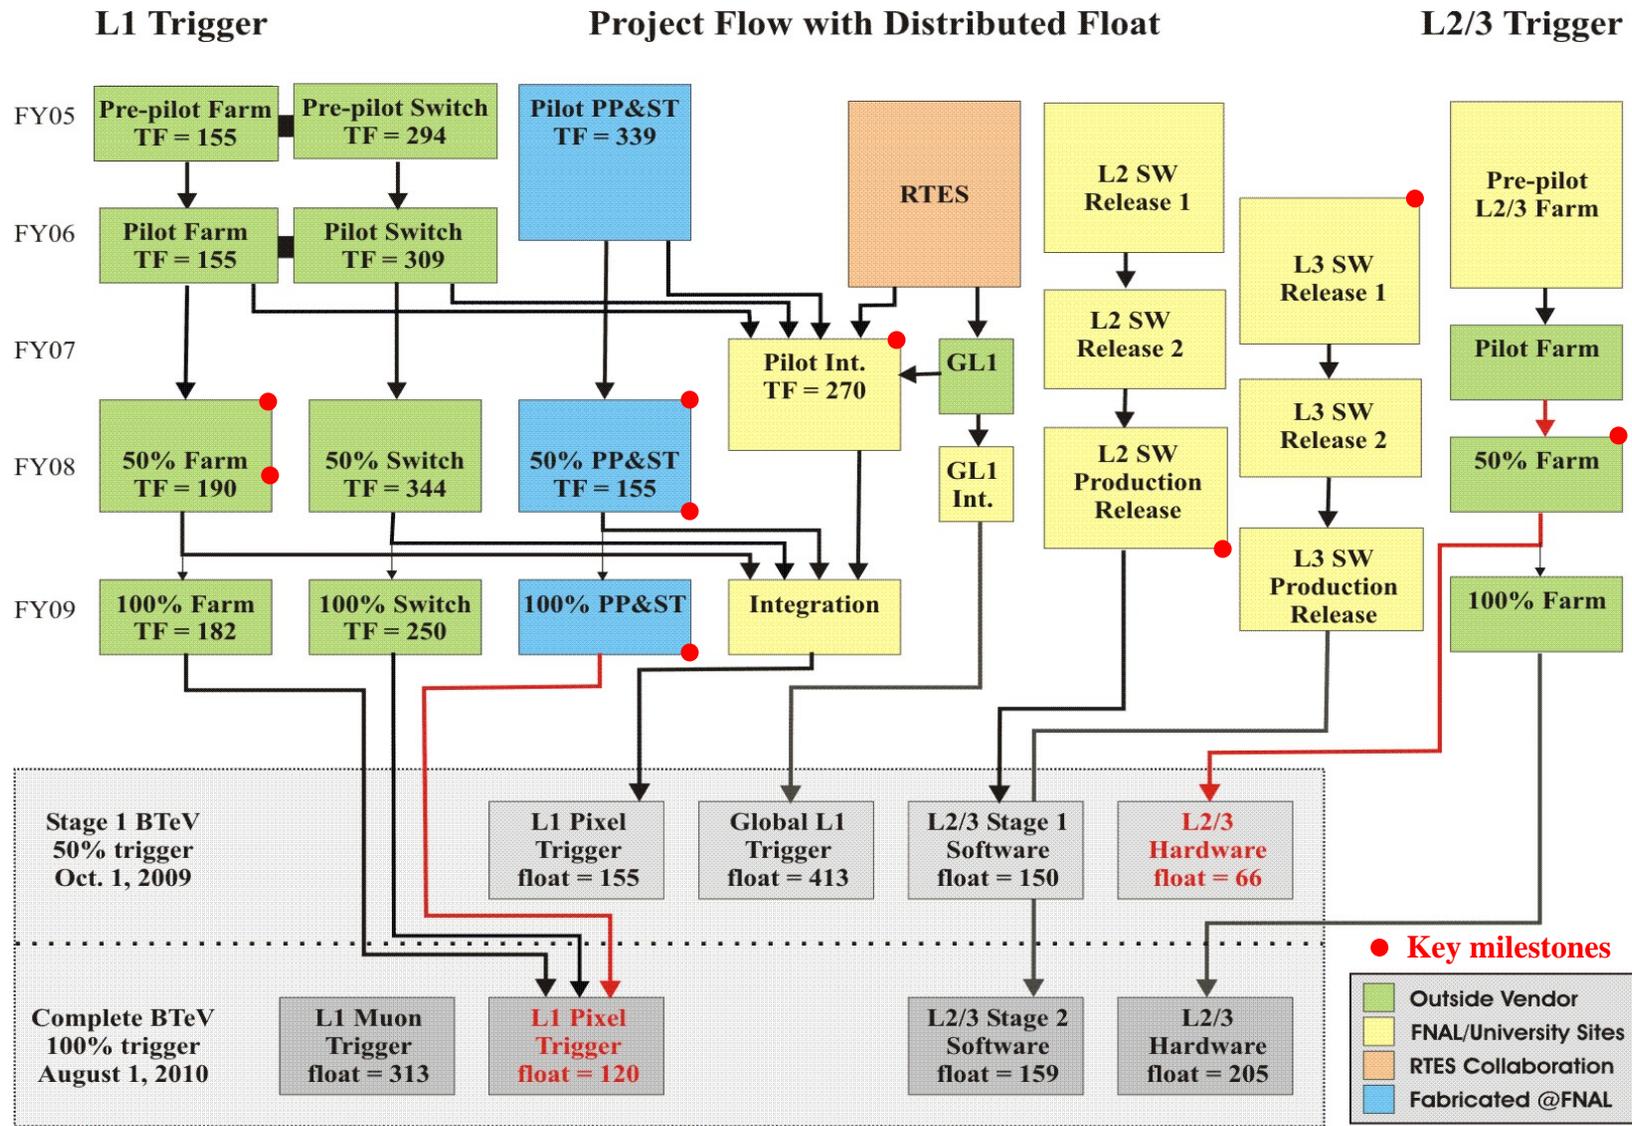
WBS 1.8

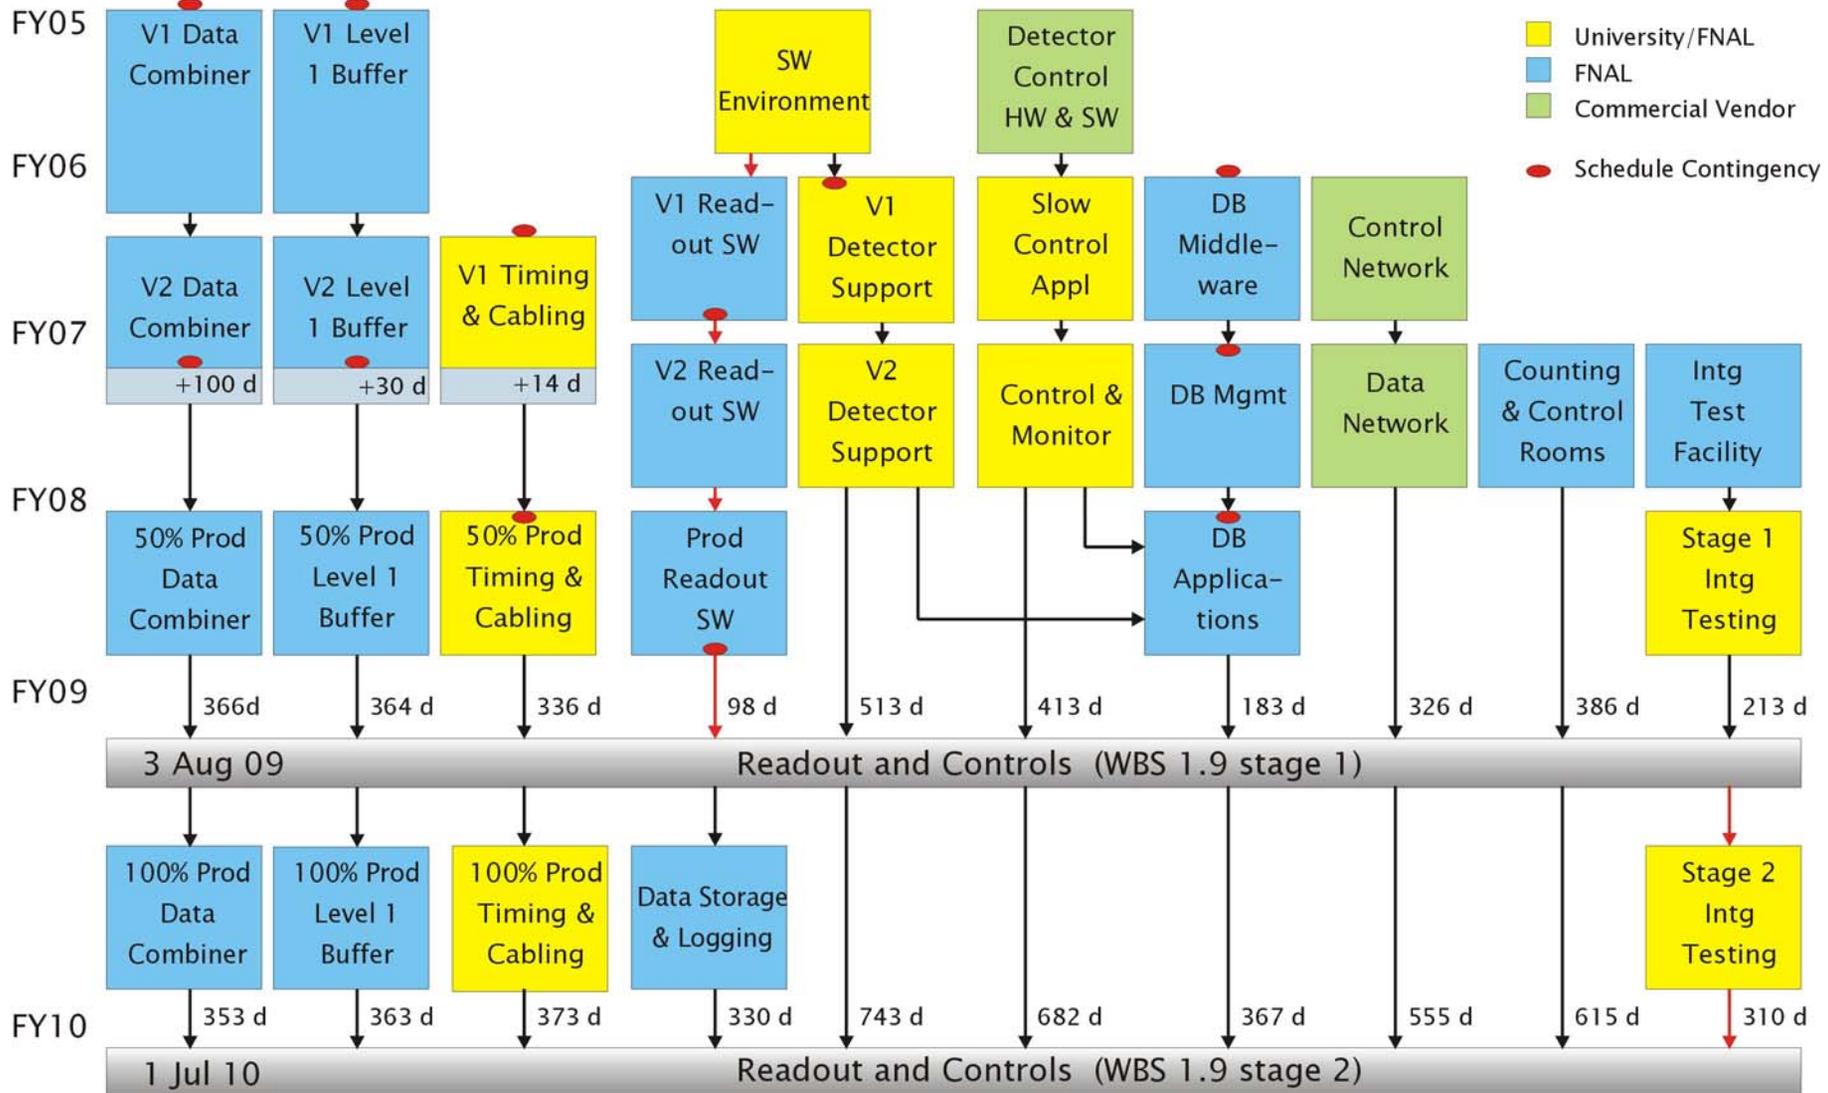
- Overview – Erik Gottschalk
- L1 pixel trigger & GL1 – Vince Pavlicek
- L1 muon trigger – Mike Haney
- L2/3 software – Paul Lebrun
- L2/3 hardware – Harry Cheung

WBS 1.9

- Overview – Margaret Votava
- Readout and controls hardware – Mark Bowden
- Readout and controls software – Margaret Votava

Additional Slides





T2	Trigger pilot system tested	Jan-07
T2,T3	Stage 1 production release of L2/3 software	Jul-08
T3	Begin L1 2-highway pixel processor & segment tracker production	Mar-07
T3	End L1 2-highway pixel processor & segment tracker production	May-08
T3	Begin L1 2-highway farm production	Mar-07
T3	End L1 2-highway farm production	Mar-08
T3	Begin L2/3 farm worker node procurement	Jul-07
T3	Begin L3 software development	Oct-05
T3	Complete trigger system and integration with DAQ	Aug-09

T2, T3	Data Combiner Board pre-production units tested and approved	Jun-06
T2	Multinode release of Data Acquisition RCS package	Nov-07
T3	Production DCB delivered and tested	Feb-09
T3	Production Level 1 Buffers delivered and tested	Jan-09
T3	Single node release of RCS package	Jun-07
T3	Data Acquisition software complete	Feb-09
T3	Calibration and Trigger database complete	Jun-08

DB	Database
DCB	Data Combiner Board
DDR	Double Data Rate
DRAM	Dynamic Random Access Memory
FPGA	Field Programmable Gate Array
GBE	Gigabit Ethernet
GL1	Global Level 1
Infiniband	Third generation high-speed networking standard
ITCH	Information Transfer Control Hardware
L1B	Level 1 Buffer
PCI	Peripheral Component Interface
PCI-Express	High-speed serial version of PCI
PP&ST	Pixel Preprocessor and Segment Tracker
PTSM	Pixel Trigger Supervisor and Monitor
RCS	Run Control System
RTES	Real-Time Embedded Systems
SODIMM	Small Outline Dual Inline Memory Module
Xserve G5	Apple's PowerPC based 1U server with dual 64-bit processors